



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

**DEFINING THE LEVELS OF ADJUSTABLE AUTONOMY:
A MEANS OF IMPROVING RESILIENCE IN AN
UNMANNED AERIAL SYSTEM**

by

Ellen M. Chang

September 2014

Thesis Advisor:
Second Reader:

Warren Vaneman
Oleg Yakimenko

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE September 2014	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE DEFINING THE LEVELS OF ADJUSTABLE AUTONOMY: A MEANS OF IMPROVING RESILIENCE IN AN UNMANNED AERIAL SYSTEM			5. FUNDING NUMBERS	
6. AUTHOR(S) Ellen M. Chang				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING /MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol number ____N/A____.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT This thesis investigates how to design in different levels of autonomy to improve the resilience of an unmanned aerial system (UAS) by applying the Function-specific Level of Autonomy Tool (FLOAT) developed by NASA. This tool helps to define the levels of autonomy human-operators are comfortable with as well as assists designers in understanding how to design in that level of autonomy. The thesis begins by reviewing past literature about resilience in engineered systems, defining terms pertaining to autonomy, introduces the concept of adjustable autonomy, and reviews the development supervisory control levels that define adjustable autonomy. It broadens the research that NASA performed and applies the tool to UAS functions. The extension of this thesis would lead to a more unified approach to defining levels of autonomy that can be adjusted for control of autonomous systems, and the development of components of software architecture that lead to greater systems resilience through integration of the human-operator in a way that is trusted. This effort is intended to create a foundation for human-centered automation to properly accommodate human-operator trust.				
14. SUBJECT TERMS autonomy, trust, resilience, adjustable, levels, FLOAT, UAS, NASA, "adjustable autonomy", "levels of autonomy"			15. NUMBER OF PAGES 91	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**DEFINING THE LEVELS OF ADJUSTABLE AUTONOMY: A MEANS
OF IMPROVING RESILIENCE IN AN UNMANNED AERIAL SYSTEM**

Ellen M. Chang
Commander, United States Navy Reserve
M.S., Naval Postgraduate School, 2014

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN SYSTEMS ENGINEERING

from the

**NAVAL POSTGRADUATE SCHOOL
September 2014**

Author: Ellen M. Chang

Approved by: Warren Vaneman
Thesis Advisor

Oleg Yakimenko
Second Reader

Cliff Whitcomb
Chair, Department of Systems Engineering

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

This thesis investigates how to design in different levels of autonomy to improve the resilience of an unmanned aerial system (UAS) by applying the Function-specific Level of Autonomy Tool (FLOAAT) developed by NASA. This tool helps to define the levels of autonomy human-operators are comfortable with as well as assists designers in understanding how to design in that level of autonomy. The thesis begins by reviewing past literature about resilience in engineered systems, defining terms pertaining to autonomy, introduces the concept of adjustable autonomy, and reviews the development supervisory control levels that define adjustable autonomy. It broadens the research that NASA performed and applies the tool to UAS functions. The extension of this thesis would lead to a more unified approach to defining levels of autonomy that can be adjusted for control of autonomous systems, and the development of components of software architecture that lead to greater systems resilience through integration of the human-operator in a way that is trusted. This effort is intended to create a foundation for human-centered automation to accommodate human-operator trust properly.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

I.	INTRODUCTION AND PROBLEM FORMULATION	1
II.	BACKGROUND AND LITERATURE REVIEW	5
A.	LITERATURE REVIEW	5
1.	Resilience.....	5
2.	Autonomy.....	10
3.	Supervisory Control.....	12
III.	METHODOLOGY FOR DETERMINING THE LEVEL OF AUTONOMY TO DESIGN INTO AN AUTONOMOUS SYSTEM	17
A.	DERIVATION OF NASA’S LEVEL OF AUTONOMY RATING TOOL “FLOAAT”	18
B.	ADAPTATION OF FLOAAT FOR A UAS	21
1.	Functional Architecture of an Unmanned Aerial System ...	21
2.	Adaptation of the Supervisory Control Reference Scale....	26
3.	Inclusion of Architectural Attributes for a Resilient System.....	30
IV.	APPLICATION OF THE FUNCTION-SPECIFIC LEVEL OF AUTONOMY AND AUTOMATION TOOL TO AN UNMANNED AIR SYSTEM.....	33
A.	CATEGORIZING UAS FUNCTIONS INTO THE OODA FRAMEWORK.....	33
B.	ASSESSING UAS FUNCTIONS USING THE FLOAAT QUESTIONNAIRE	36
C.	REVIEWING THE RESULTS IN SUMMARY.....	38
D.	SUMMARY AND LESSONS LEARNED	42
V.	FLOAAT APPLICABILITY ASSESSMENT AND FUTURE RESEARCH.....	45
A.	EVALUATION OF ADAPTING FLOAAT TO A UAS	45
1.	The FLOAAT Questionnaire	46
2.	Employment of OODA Categories	47
3.	Employment of the 10-Level Rating Scale	47
4.	Incorporation of Architecting for Resilience	48
B.	ADDITIONAL CONSIDERATIONS AND FUTURE RESEARCH.....	48
1.	Human-Operator Collaboration.....	49
2.	Verification, Validation, and Certification	49
3.	Autonomy Architecture.....	49
	APPENDIX A. LEVEL OF AUTONOMY QUESTIONNAIRE	51
	APPENDIX B: EXAMPLE OF NASA’S LEVEL OF AUTONOMY SCORES FOR A SPACECRAFT	59
	BIBLIOGRAPHY	61
	INITIAL DISTRIBUTION LIST	69

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1.	Architectural attributes that enable a resilient systems architecture (from Vaneman and Triantis 2014)	7
Figure 2.	List of Resilience Heuristics (from Madni and Jackson 2009).....	8
Figure 3.	Levels of Autonomy from a Computer's Perspective (from Parasuraman 2000)	14
Figure 4.	Contextual Use Cases for an Unmanned Control System (from UCS Control Segment Working Group 2014)	22
Figure 5.	Functional Architecture of a UAS	23
Figure 6.	Functions Performed by UAS Segments	23
Figure 7.	Cooper-Harper Handling Qualities Rating Scale (from NASA History Web Site 2014b)	27

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

Table 1.	Sheridan Model of Autonomy (from Sheridan 1992)	13
Table 2.	Description of FLOAAT Tool Questions (from Proud and Hart 2005) .	20
Table 3.	A Set of UAS Mission Functions (from UCS Working Group 2014)	25
Table 4.	NASA's 5-pt Scale and Definition of the Levels of Autonomy (from Proud and Hart 2005)	28
Table 5.	Level of Autonomy Reference Rating Scale for a UAS (from Proud and Hart, 2005)	29
Table 6.	Adaptation of FLOAAT to Consider Architectural Attributes of Resilient Systems (from Proud and Hart, 2005)	31
Table 7.	UAS Functions defined and categorized	35
Table 8.	Questions from the FLOAAT Questionnaire and Comments to Those Questions for the Weather and Environmental Data Function (adapted from Proud and Hart, 2005)	37
Table 9.	Weather Function (from Proud and Hart, 2005)	38
Table 10.	Level of Autonomy Scores for UAS Functions (after Proud and Hard 2005)	40
Table 11.	Summary Levels of Supervisory Control for Adjustable Autonomy (from Sheridan 1992 and Parasuraman 2000)	41
Table 12.	Level of Autonomy Questionnaire (from Proud and Hart, 2005)	51
Table 13.	NASA's FLOAAT Scores for a Spacecraft Re-planning tool (from Proud and Hart 2005)	59

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

DOD	Department of Defense
DSB	Defense Science Board
FLOAAT	Function-specific Level of Autonomy Assessment Tool
ISR	intelligence, surveillance, and reconnaissance
LoA	Level of Autonomy
NASA	National Aeronautics and Space Administration
NACA	National Advisory Committee for Aeronautics
OODA	observe, orient, decide and act
OSD	Office of the Secretary of Defense
UAS	unmanned aerial system
UAV	unmanned aerial vehicle
UCS	UAS control system

THIS PAGE INTENTIONALLY LEFT BLANK

EXECUTIVE SUMMARY

In 2012, the Defense Science Board published a study on autonomy. This comprehensive look at autonomy and its employment in the DOD provided many areas of improvement, one being the interface between the operator and autonomous systems. According to the Defense Science Board's 2012 Task Force Report: *The Role of Autonomy in DoD Systems*, the interface between the unmanned system and operator is brittle; this brittleness was noted as a limitation preventing further adoption of autonomous systems (DSB, 2012). Brittleness, in this case, was where these autonomous systems, being too deterministic, were not able to adapt to situations that were different from anticipated. Designers are not able to anticipate fully all of the scenarios that could arise during the use of the system. Hence, the research focus on building a system that can react and adapt—to design and build in robustness to counteract the unintended brittleness, and to leverage the human-operator in the process.

This thesis explores how to build in resiliency by providing the human-operator different levels of control over an autonomous system—ranging from fully manual to fully autonomous. It does so by adapting the Function-specific Level of Autonomy Tool (FLOAAT) developed by NASA for application on a UAS. By properly defining and designing in to the system different levels of autonomy that the human-operator can select, it improves human-system interaction in a way that optimizes each the competencies of both the human operator and the system.

In summary, FLOAAT proved to be an effective tool to get at the heart of what level of autonomy is appropriate for any one function. The approach forced thoughtful consideration of different design, employment and cost aspects of making a function autonomous, which, in a manner, forced thorough requirements analysis for that function. Employment of FLOAAT showed that the process for determining the Level of Autonomy for any one function is iterative; a subject matter expert, in working through the questions and rating level

definitions wrestles with the derived level resulting from the tool, and, conceivable would negotiate the intent and meaning of this level with a broader systems design team.

Though the tool has proven useful in initial research, further investigation is required to truly validate its employment in the UAS domain. NASA has applied this to several programs. In doing so, they have developed approaches to validate the level of autonomy as suggested by FLOAAT (Proud and Hart 2005). They have a baseline of experience to draw from. This is not the case for unmanned aerial systems. If this tool were to be more widely adopted, there is more work to be done:

1. Determination of the composition of the team who should participate in the process of defining the level of autonomy by answering the questionnaire. How many and of what type of subject matter experts should be included?
2. According to Proud and Hart, the Level of Autonomy tool employed and adapted was originally designed to ascertain the division of labor between the computer and the human-operator. Additional questions could be added to determine the division of labor between what should be on the aircraft and what functions should be executed in the mission control system.
3. The questions should be refined and tested against a larger cross-section of users or subject matter experts to ensure the question is clear and the intent is communicated.
4. Test cases should be developed in order to more quickly validate the scores and even prototype the output.

In conclusion, autonomous systems have been changing the way the military does business, and, with recent investment by the DOD and the commercial world, is on the threshold of exerting deep changes in military operations. These systems can and will be able to be operated without direct human control for extended periods of time and over long distances. This is beneficial and will open the field for more applications while reducing costs; but, it should be done with the human-operator and his/her strengths and weaknesses, in mind. Or else, the systems may not be adopted, or, even worse, the systems

may not be safe. As such, the following are a few suggested areas of further research:

1. Human-Operator Collaboration
 - Determine how the roles of human-operations and the autonomous systems, as well as the human-system interface, should evolve to enhance more efficient yet safe operations.
 - Further understanding of human psychology in the operation of autonomous systems.
 - Interfaces, be they visual, aural, focused on assistance or alerting to problems that improve human performance.
 - Approaches to adjust to different skill and cognition levels in human-operators, with an eye toward safety.
2. Verification, Validation, and Certification
 - Develop standards and processes for the verification, validation, and certification of autonomous systems, and determine their implications for architecture and design.
3. Autonomy Architecture
 - Explore and define the landscape of autonomous systems architecture to further the ability to adapt and verify and validate the system.

LIST OF REFERENCES

- Defense Science Board (DSB). 2012. *Task Force Report: The Role of Autonomy in DoD Systems*. Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, July 2012.
<http://www.acq.osd.mil/dsb/reports/AutonomyReport.pdf>.
- Proud, Ryan W., and Jeremy J. Hart. 2005. "Function-specific Level of Autonomy and Automation Tool (FLOAAT) Rendezvous, Proximity Operations, and Docking (RPOD) Reference Levels of Autonomy and Automation" (NASA/JSC Document Number AFMFLOAAT002, Output Version). Unpublished.

THIS PAGE INTENTIONALLY LEFT BLANK

ACKNOWLEDGMENTS

To my husband, Doug, who supported me through and through!

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION AND PROBLEM FORMULATION

This thesis examines an approach to improve the resilience of autonomous systems using adjustable autonomy and namely focuses on an unmanned aerial system. It leverages a framework the National Aeronautics and Space Administration (NASA) has employed in its space systems to enhance autonomy (Proud and Hart, 2005) while ensuring the trust of the operator. In doing so, the method developed provides a means to design in levels of autonomy that engenders the trust of a human-operator and that provides a means to improve the resiliency of autonomous systems.

The motivation for this focus came from the examination of the Defense Science Board's (DSB) report on autonomy completed in 2012 (Defense Science Board [DSB] 2012). This comprehensive look at autonomy and its employment in the DOD identified many areas of improvement, one being the interface between the operator and autonomous systems. The interface between the unmanned system and operator was characterized as brittle; this brittleness was noted as a limitation preventing further adoption of autonomous systems (DSB 2012). Brittleness, pointed out in the DSB report, was where these autonomous systems, by being too deterministic, were not able to adapt to situations that were different than anticipated when the software was originally developed. In essence, this brittleness is the opposite of resiliency. Resiliency is the ability to adapt to changing conditions (natural or man-made) through planning on how to absorb (withstand) and rapidly recover from adverse events and disruptions (Vaneman and Triantis 2014).

Brittleness can arise because of the inability to anticipate and design for all of the scenarios that could arise during the use of the system (Duda and Shortliffe 1983). While this definition covers system operations in predictable environments, it breaks down in the context of uncertainty (Lenat and Guha 1989). Failure modes cannot be exhaustively anticipated; designers already are challenged to think through as many scenarios as possible. Hence, the focus on

building a system that can react and adapt—to design and build in robustness to counteract the unintended brittleness.

An approach to engineering resilience into a system is to leverage the capabilities of the human operator. Humans may not be able to land a plane in a predefined precise location as well as a computer can, but humans can adeptly and much more effectively anticipate issues or unforeseen events and adjust their response toward mission achievement. To be effective, autonomous systems need to be competent collaborators with human-operators. Critical analysis to define the appropriate functional allocation of the roles between the system and human, and level of automation to those functions, are essential. A compromise has to be found between completely manual and fully autonomous operations. This is where adjustable autonomy comes in, where control is provided to the human-operator to enable a level of autonomy with which the operator is comfortable. Such interaction allows the dynamic adjustment of autonomy to face whatever situation or environment exists at that time (Zieba et al. 2009).

This thesis leverages a tool developed by NASA that assists with determining what level of autonomy to design, function by function, into a system. The approach involved adapting the Function-specific Level of Autonomy and Automation Tool (FLOAAT), by considering supervisory-control principles developed for air systems and architectural attributes for resilient systems as defined by Vaneman and Triantis (Vaneman and Triantis, 2014). Chapter II provides the background and summarizes research pertinent to the domains of engineering resilience into a system, autonomy and supervisory control. Chapter III breaks down NASA's approach to adjustable autonomy and illustrates how the tool was adapted and why. Chapter IV summarizes and presents the application of NASA's FLOAAT to an unmanned aerial system. Chapter V then reviews and summarizes conclusions, discusses lessons learned, and suggests research that is required to advance the domain. Included are two appendices: Appendix A, which provides details of the 35 questions posed and Appendix B, which

provides an excerpt of NASA's scoring for a set of functions as a foundational reference.

The key questions posed in this thesis are:

1. How can one design in proper levels of autonomy to optimize the human-operator team?
2. How can one design in levels of autonomy to enable greater systems resilience?
3. What aspects of the derived level of autonomy and design information can be modeled in order to test the autonomous system?
4. How can one "architect" resilience into autonomous systems in order to enhance the manned-unmanned interactions, engender trust, and reduce instead of increase the workload?

The benefits of this research and study include:

1. Providing an initial foundation of research into defining levels of autonomy and assess benefits of furthering such research.
2. Lessons learned regarding the process of defining levels of autonomy and whether or not NASA's Function-specific Level of Autonomy and Automation tool has merit as applied to an unmanned aerial system.
3. Lessons learned about the definition of adjustable autonomy as it applies to architecting a resilient system, leveraging the framework provided by Vaneman and Triantis.
4. Recommendations on designing the human-system interface in order to engender trust and allow for advancement in the employment of autonomous systems.

THIS PAGE INTENTIONALLY LEFT BLANK

II. BACKGROUND AND LITERATURE REVIEW

Resilience connotes the ability to spring back, to recover. Today, with the advancement of autonomy and robotics, attention has also moved toward applying the concept of resilience to these engineered systems where reliability and fault prediction or failure modes and effects have dominated system design. When examining autonomy in 2012, the Defense Science Board noted that systems were brittle as opposed to resilient (DSB 2012). Brittleness arises when a system is not designed to work in all of the scenarios that could be encountered during the use of the system. This is of concern for deterministic systems, as failure modes cannot be exhaustively anticipated. Hence, the focus of this research is on methods to build a system that can react and adapt to counteract the unintended brittleness.

The Introduction of this thesis sets the stage by providing an overview of the definition of resilience as it applies to engineered systems and systems of systems. It also suggests that the field of adjustable autonomy could be an approach improves system resilience. This chapter takes the next step and provides more detailed discussion of relevant technical terms and presents an overview of research in the domains of resilience and autonomy. It also provides an overview of a tool that NASA developed to define and design autonomy levels. This thesis proposes to take this tool and adapts it for employment on an unmanned aerial system.

A. LITERATURE REVIEW

1. Resilience

Within technical fields, the use of the term resilience has a tradition in materials science (Martin-Breen and Anderies 2011). Martin-Breen and Anderies suggest that the definition of resilience should be customized to the discipline to which it is applied. Only then is context, which can be important, considered. They also note the problem with systems is that they are deterministic; they

cannot adapt. Hence, Martin-Breen and Anderies suggest there is a need to design in the means for systems to adapt in order to be resilient (Martin-Breen and Anderies 2011).

Vaneman and Triantis studied resilience engineering in a system of systems context and have proposed definitions appropriate for this domain. In their presentation they note that Resiliency is the ability to adapt to changing conditions (natural or man-made) through planning on how to absorb (withstand) and rapidly recover from adverse events and disruptions (Vaneman and Triantis 2014). Important to systems design and engineering efforts, Vaneman and Triantis also address resilience in systems architecture. They note that systems architecture is resilient if it can provide the necessary operational functions, with a higher probability of success and shorter periods of reduced capabilities before, during and after an adverse condition or disruption through avoidance, robustness, recovery and reconstitution (Vaneman and Triantis 2014). In doing so, they suggest four architectural principles to strive for:

- Avoidance: proactive or reactive measures taken to reduce the likelihood or impact of adverse conditions or threats.
- Robustness: design feature to resist functional degradation and enhance survivability.
- Recovery: actions and design features that restore a lost capability to meet a specific mission set (perhaps the most critical mission set).
- Reconstitution: actions and design [that] features a measure of how much the total capability can be replaced, and the time it takes to achieve [the replacement] (Vaneman and Triantis 2014).

These elements are further defined by architectural attributes as listed in Figure 1 below. The figure shows, as an example, that the ability for a system to avoid degradation comes from operational flexibility, flexibility in policies and procedures, loose coupling, and extendibility. Vaneman and Triantis suggest a set of attributes for each of the architectural principles and imply that

consideration of them early in the lifecycle will aid in designing in resiliency into a system (Vaneman and Triantis 2014).

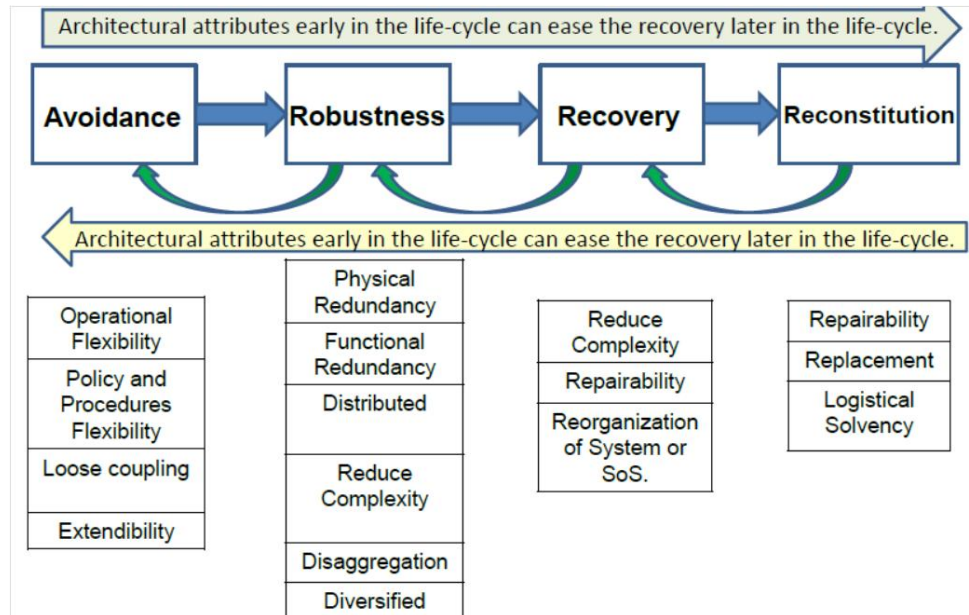


Figure 1. Architectural attributes that enable a resilient systems architecture (from Vaneman and Triantis 2014)

Consideration of these elements help frame the design process to consider these attributes early in the architecting process. As an example, thinking through how to enable “policy and procedures flexibility” early on may result in derived requirements that provide the option for a higher number of levels of autonomy in order to provide a greater degree of flexibility. In fact, consideration of this and other architectural attributes that comprise “Avoidance” has influenced the definition of adjustable autonomy for an Unmanned Aerial System (UAS), discussed in the next chapter. This design consideration needs to be contemplated up front. If they were not, the development of the system could incur significant additional cost if the applications were to be redesigned to improve “policy and procedure flexibility”. It is this set of architectural attributes that are later employed to help designers consider elements that improve resilience.

In a fairly recent work, “Towards a Conceptual Framework for Resilience Engineering,” Madni and Jackson provide a framework to look at engineering in resilience (Madni and Jackson, 2009). Their section on heuristics describes fourteen attributes that characterize resilient systems. Figure 2 below lists these fourteen heuristics.

Functional Redundancy	there should be alternative ways to perform a particular function that does not rely on the same physical systems
Physical Redundancy	employ redundant hardware (e.g. processors) to protect against hardware failure when functional redundancy is not possible
Reorganization	system should be able to restructure itself in response to external change
Human Backup	humans should be able to back up automation when there is a context change that automation is not sensitive to and when there is sufficient time for human intervention
Human-in-the-Loop	humans should be in the loop when there is a need for "rapid cognition" and creative opinion generation
Predictability	automated systems should behave in predictable ways to assure trust and not evoke frequent human over-ride
Complexity Avoidance	systems should reflect system complexity and not complexity added by poor human design practices
Context Spanning	system should be able to survive most likely and worst case scenarios, either natural or man-made
Graceful Degradation	systems performance should degrade gradually when the unexpected occurs for which system is not prepared
Drift correction	system should be able to monitor and correct drift toward brittleness by making appropriate tradeoffs and taking timely preventative action
"Neutral" state	system should be able to prevent further damage from occurring when hit with an unknown perturbation until problem can be diagnosed
Inspectability	system should allow for human intervention needed without requiring humans to make unsubstantiated assumptions
Intent Awareness	system and humans should maintain a shared intent model to back up each other when called upon
Learning/ Adaptation	continually acquiring new knowledge from the environment to reconfigure, reoptimize and grow

Figure 2. List of Resilience Heuristics (from Madni and Jackson 2009)

Heuristics are experience-based frames of reference to employ when thinking about a topic. Of the fourteen, two, highlighted in yellow in the figure, consider that humans are *essential* elements of a resilient system. Humans should be able to backup automation when change occurs and that humans should be in the loop when there is a need for rapid cognition (Madni and Jackson 2009).

Jackson and Ferris (2012) take these concepts further when they establish the “*human in control*” principle. They posit that the human operator should retain final decision making authority in critical situations unless the pressure of time demands a quick decision. They further list an “*automated function*” principle that suggests automating certain types of tasks: (1) those that need to be performed quickly, in a split second; (2) those that are not too complex—where their definition of complexity does not include an uncertain, unpredictable situation; and (3) where a task is boring, repetitive, or distracting. (Jackson and Ferris 2012). The latter two address the role of a human and the role of a robot, respectively. These two principles are important to adjustable autonomy as they provide frames of reference with which to consider different tasks and functions and how automated they should be.

It is important to highlight these principles in the context of resilience as it reinforces that human-operators are important components of a system and should be positioned appropriately in any system in order to improve resilience of that system, especially when included in aspects or during times where automation backup is required, when the human-operators anticipatory skills are needed. Conversely, the principles note where human-operators are NOT ideal—when tasks are extremely fast, boring and lengthy, or repetitive.

Zieba et al. (2009) formally link adjustable autonomy with resilience. They note that adjustable autonomy is a means to adapt the system to situations anticipated and unanticipated. They show how a system’s ability to recover is enhanced when the human performs what he is best at and the robot performs what it is best at. The results of their experiment illustrate how the human-

operator and system, as a collective, work together to achieve the mission at a higher level of resilience, based on measures they have assigned. When properly designed and positioned, they conclude that human-robot collaborative control is a means to increase the resilience of an autonomous system. (Zieba et al. 2009). Hence, it is important to design in autonomy to optimize the division of control between the human and the system to result in a more resilient system.

2. Autonomy

An autonomous system has to be resilient in order to adapt to unplanned events. Resilience heuristics provide a framework to employ in autonomous systems design, enabling systems to adapt and recover from unanticipated problems. This activity, in and of itself, improves resilience of a system. But, before discussing autonomy and why adjustable autonomy improves the resilience of an unmanned system, it is important to define the terms and provide background on research that has shown how the two relate.

Autonomy does not have a single unified definition. In fact, the word is more widely used in social, political and psychological domains, where it connotes self-determination (Christman 2009). The autonomous systems domain that has evolved since the 1950s and 1960s has adopted the word for the simple reason that it does mean to self-determine how to operate. Gregory Dorais, a NASA expert and researcher on intelligent systems and human-centered automation, is considered as one of the pioneers of “adjustable autonomy.” His research dates back to the early 1990s. He defined an adjustable autonomous system as a control system that has the ability to 1) be completely in control; 2) be able to supervise manual control; or 3) be able to shift among these control extremes in a safe and efficient manner (Dorais and Kortenkamp 2008). Further, a system’s “adjustable autonomy” can involve changes in: the complexity of the commands it executes; the resources (including time) consumed by its operation; the circumstances for when it will request user information or control; the circumstances when it will override or allow manual control; the number of

subsystems that are being controlled autonomously (Dorais and Kortenkamp 2008).

Chad Frost, then director of the NASA Ames Autonomous Systems and Robotics Division, clarified in a speech he gave in 2010 the difference between autonomy and automation. He noted that “Many definitions are possible...but here we focus on the need to make choices...an automated system doesn’t make choices for itself, it follows a script. If it encounters an unplanned situation, it stops and waits for human assistance. Thus, for an automated system, choices have either already been made or they must be made externally. By contrast, an autonomous system does make choices on its own and tries to accomplish objectives without human intervention, even when encountering uncertainty” (Frost 2010, pg 2).

Relevant research in autonomous systems can be classified under five areas: autonomous robots, tele-operation, adjustable autonomy, mixed initiatives, and advanced interfaces (Goodrich et al. 2001). An autonomous robot is a robot that performs behaviors or tasks with a high degree of autonomy; an autonomous robot may also learn or gain new knowledge like adjusting for new methods of accomplishing its tasks or adapting to changing surroundings (Dorais and Kortenkamp, 2008).

Telerobotics is the area of robotics concerned with the control of semi-autonomous robots from a distance, using wireless or tethered connections (Sheridan 1992). Mixed-initiative systems integrate human and automated reasoning to take advantage of their complementary reasoning styles and computational strengths (Tecuci et al. 2003). This area of research addresses the division of responsibility between the human and autonomous system, control, shared awareness, exchange of information/knowledge, and situation evaluation.

The adjustable autonomy domain takes into account the adjustment of the degree or level of autonomy a system exhibits. It keeps the human-operator

involved, in control, by allowing him to trade-off between the convenience offered by autonomy, and the amount of control he would like to exert (Dorais and Kortenkamp, 2008). When humans and machines share responsibility for achieving a specific task, responsibility can be thought of as shifting between the human and the robot along a continuum of fully manual or fully automated. This underpins the concept of adjustable autonomy. Adjustable autonomy was introduced for supervisory control of robotic systems at specific levels along this continuum (Bonasso et al. 1997). NASA notes that much of their approach derives from Bonasso et al. (1997), where manual and automated control methods for each task was defined and the level of autonomy could be selected.

3. Supervisory Control

Given that the goal is to design in a level of autonomy—or a form of “supervisory control,” it is necessary to discuss research in this domain. The term “supervisory control” is a general term for control of a control system. A control system is a device, or set of devices, that manages commands, directs or regulates the behavior of other device(s) or system(s) control loops, whether by a human or an automatic control system. When contemplating autonomy, there naturally arises the question of the level of control that a human-operator should have or desires.

Sheridan (1976) was one of the first to assign control levels to robots. Most autonomy researchers use this as a reference for an initial understanding of how humans and computers interact. Sheridan focused on telerobotics where the human is physically separated from the system, but still issuing commands (Sheridan and Johannesen 1976). The most relevant information comes from Sheridan’s work on trust development, such as reliability, robustness, familiarity, usefulness, and dependence (Sheridan 1992). Note some of these trust attributes are the same or similar to those architectural attributes called for in architecting resilient systems. In dealing with these trust issues, Sheridan proposed a ten-level scale of automation as depicted in Table 1 (1992).

Examination of the levels note that Levels 2 through 4 are centered on who makes the decisions, the human or the computer. Levels 5–9 are centered on how to execute that decision. Levels 1 and 10 are end-bounds for either extreme (Sheridan 1992).

Table 1. Sheridan Model of Autonomy (from Sheridan 1992)

“Sheridan” Model - levels of autonomy
1) Computer offers no assistance, human must do it all.
2) Computer offers a complete set of action alternatives, and
3) narrows the selection down to a few, or
4) suggests one, and
5) executes that suggestion if the human approves, or
6) allows the human a restricted time to veto before automatic execution, or
7) executes automatically, then necessarily informs the human, or
8) informs him after execution only if he asks, or
9) informs him after execution if it, the computer, decides to.
10) Computer decides everything and acts autonomously, ignoring the human.

Subsequently, in 2000, Parasuraman provided a revised model for the levels of automation. His model kept the ten levels but split the tasks performed into four categories: “information acquisition; information analysis; decision and action selection; and action implementation” (Parasuraman 2000). Parasuraman’s framework is depicted in Figure 3 below.

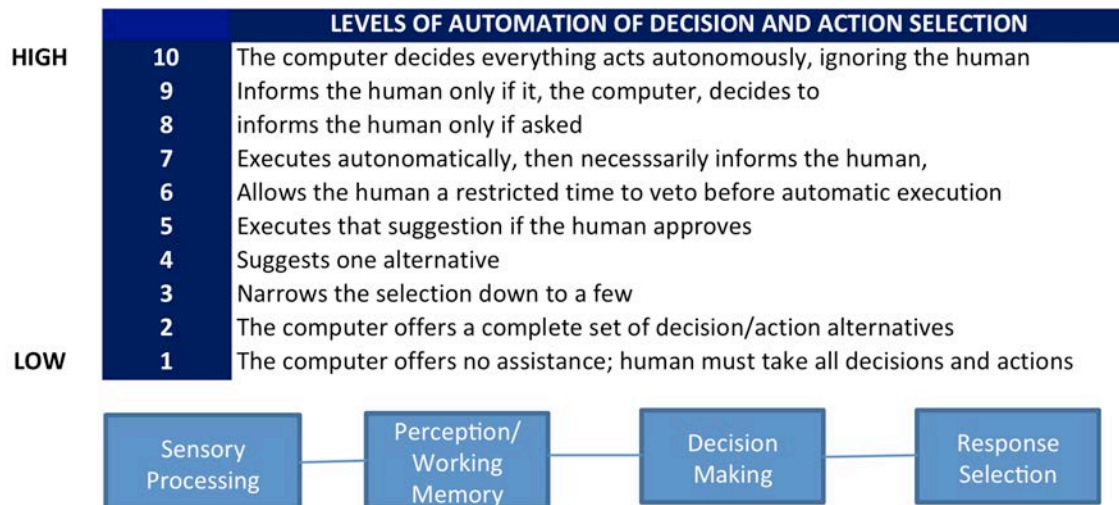


Figure 3. Levels of Autonomy from a Computer's Perspective
(from Parasuraman 2000)

Information acquisition is the task of sensing, monitoring, and bringing information to a human's attention (Parasuraman, 2000). Information analysis is performing all of the processing, predictions, and general analysis tasks. Decision and action selection result in making choices. For example, "Based on the available analysis, what should the system do?" Action implementation is acting on decisions or commanding new actions. Levels in this category include the computer asking for authority to proceed and allowing human overrides. The breakdown of a decision in this decision making sequence enabled more precise interpretation of a level of autonomy for any one particular system (Parasuraman 2000).

Parasuraman's 4 categories approached a task from the perspective of the computer. NASA's Proud and Hart (2005) switched the perspective to that of a human-operator and took a chapter from military decision-making. They saw a similar 4-tiered system in Boyd's Observe, Orient, Decide and Act (OODA) framework (Boyd 1996). According to Proud and Hart (2005), Boyd's system added two important characteristics--feedback and implicit control. Feedback is obtained during the decision cycle, and decisions do not necessarily have to become actions. Decisions themselves can spark new analysis tasks or requests

for new observations, but not result in actions (Proud and Hart 2005). And, implicit control refers to the fact that there could be an entity that retains control whether or not there is an explicit action. The example given is a flight control processor that has implicit control of the vehicle and will continue to carry out the management of the vehicle unless otherwise commanded (Proud and Hart 2005).

Since this thesis involves assessing levels of autonomy of a UAS, it is important to present research relevant to supervisory control in the aviation domain. In the world of aircraft control, the first control rating was developed by the National Advisory Committee for Aeronautics (NACA), the predecessor to NASA, called the Cooper-Harper rating scale (Lintern and Hughes 2008; NASA History Web Site 2014a). The Cooper-Harper rating scale is a set of criteria used by test pilots and flight test engineers to evaluate the handling qualities of aircraft during flight test. The scale ranges from 1 to 10, with 1 indicating the best handling characteristics and 10 the worst. (Lintern and Hughes 2008). The Air Force Research Lab has since investigated several permutations of the Cooper-Harper scale, and others have developed alternatives, in attempts to overcome some of the flaws of the Cooper-Harper scale. However, it is the one scale that is consistently employed. Further reference and linkage will be discussed in Chapter III.

THIS PAGE INTENTIONALLY LEFT BLANK

III. METHODOLOGY FOR DETERMINING THE LEVEL OF AUTONOMY TO DESIGN INTO AN AUTONOMOUS SYSTEM

Increased autonomy in unmanned systems, trusted by the operators, is necessary for the next-generation of Naval Airborne systems, to meet cost, safety, and mission requirements. Detailed, manual work sometimes makes humans feel in control; in reality, humans may not be the most suitable for some of the functions they perform. But humans will not willingly relinquish control unless they are confident that what they are giving control over to will do the job just as well. They need to trust the system. Ideally, the autonomous system should augment the human, perhaps raising their capability to a higher level, and not just take over certain tasks. An example is the F/A-18 Launch function; it is autonomous because at low speeds the F/A-18 is highly unstable and human control is too slow to operate it successfully (Isby 1997). The human is there, ready to fly the plane, but is not in control of this specific function where his reaction time is too long.

Building on the foundation provided by in the literature review, this chapter details the methodology applied to modify NASA's Function-specific Level of Autonomy Tool to an unmanned system. Why NASA and why this tool? For one thing, NASA has pioneered the development of autonomous systems. Launch of Deep Space 1 almost fifteen years ago, in 1998, demonstrated the feasibility of a fully autonomous spacecraft (Frost 2010). Much of the research on autonomous systems and their interface with humans have been conducted by NASA as they embarked on this journey to enable autonomous space missions. Furthermore, this tool has some pedigree, as it has been applied and used on several of the systems, to include the Centaur, the International Space Station, the Crew Entry Vehicle and others (Proud and Hart 2005). Therefore, it has foundation and has been verified and validated in several different applications. Thirdly, it directly addressed trust issues that humans have with autonomous systems by limiting the amount of autonomy to be made available to that amount which the

questionnaire-tool noted as the notional limitation in trust. And, finally, the tool employed autonomy level definitions that were meant for designers of the system. This is an important distinction and frame of reference. Levels of autonomy are necessary for researchers to partition and frame the problem, for designers to engineer the system, and for operators, to define the control they trust.

A. DERIVATION OF NASA'S LEVEL OF AUTONOMY RATING TOOL "FLOAAT"

The space community had been investigating how to properly design in autonomy since the early 1990s (Dorais et al. 1998). The community started to implement an increasing amount of autonomy, looking at automating functions that had been manually managed by human operators in ground control. This initial foray into the field was motivated by potential cost savings (Proud and Hart 2005). Computer advancements, the emergence of highly reliable decision-making algorithms, and the emphasis on efficiency made this possible. However, they discovered that for some human spaceflight applications, full autonomy was impractical (Dorais et al, 1999). What NASA learned was that a balance had to be struck between how much human operators trusted automation, and how much benefit and cost savings automation provided (Proud and Hart, 2005).

NASA's motivation was cost; its issue was trust (Proud and Hart 2005). In the early 2000s, Proud and Hart embarked on defining an engineering approach to 1) define levels of autonomy that incorporated the concept of trust into functions and 2) enable the definitions to be used as systems requirements (Proud and Hart 2005). NASA scientists identified system functions and analyzed each function to determine how much autonomy could be tolerated and what level of autonomy to design for that function (Proud and Hart 2005). They ascertained the level by implementing a tool—a questionnaire that elicited from a set of human operators what they thought should be automated. The employed Level of Autonomy Tool was designed to determine the division of labor between computers and humans for the various functions. Adaptations were made to add

questions that ferret out design issues addressing division of labor between the ground and onboard systems. NASA implemented its approach in a tool called the Function-specific Level of Autonomy Assessment Tool (Proud and Hart 2005). FLOAAT essentially is a mechanism to assess the levels of supervisory control desired for a specific function. The tool contains 35 questions about the execution of a select number of functions for a system. The questions are intended to draw out from subject matter experts, some being engineers of a certain domain, some being designers, some being human operators, on how confident he/she would be if a computer would have full control of the said functions. A portion of the questions also draw out whether there is a return on investment, or what the cost was for the benefit of automating a certain function. The cost could, conceivably, be too high.

NASA divided the 35 questions into two primary categories—a set that gets at the heart of how much a human-operator would trust the system if the function were to be handled fully autonomously; these are called “Trust Questions,” and a second set that considered the cost of designing and developing a capability to make a function fully autonomous; these are called “Cost/Benefit Questions.” Trust questions number 20 and address issues such as complexity, software design capacity, robustness, art vs. science and other similar trust issues. Cost questions number 15 and cover categories such as usefulness (of automation), timeliness, criticality and safety. Table 2 depicts the categories and general description of what the questions are trying to elicit. The full set of the questions for the two categories of trust and cost/benefit can be found in Appendix A. One can see that the questions span a wide range of areas that are all pertinent to getting at the heart of not only trust in autonomous systems, but also good design.

Table 2. Description of FLOAAT Tool Questions
(from Proud and Hart 2005)

Category	Nature of Questions
Trust Questions - 20	
Ability	This category attempts to derive what level of ability system designers are required to have to be able to develop the algorithm and integrate it into the function correctly, and within the timeline required.
Difficulty	How difficult/complex of a design effort is it to properly automate the function in question. Questions get at technical difficulty and schedule difficulty (i.e within the timeframe required)
Robustness	Questions in this category attempt to define how robust of a design is required for the function- is there an opportunity for an "out of the box" scenario to occur.
Experience	Questions get at what experience exists in automating the function in question. The more experience of having a certain function automated, the more general knowledge exist on how to properly design the function and how human-operators react and handle situations involving the function. How autonomous (what level) has the function been shown to perform?
Understandability	This category gets at the complexity of the function and how much understanding (to include training) does the human-operator need. Do operators have a mental model of how the function should work? Understanding pertains to not just the function, but understanding the mission environment and knowing what to do next.
Art vs Science	This category attempts to derive if the function could be performed by humans based on using their experience (Heuristics) vice a fully optimal solution
Familiarity	This category derives information on if the human-operators would be familiar with output of an agent or if the function were fully automated. "How natural would the output feel"?
Training	What is the probability that the computer could come up with an answer that is "more accurate" than a human?
Override	This category derives whether or not an override is necessary - which it usually is. One of the issues is V&V of a fully autonomous agent performing certain functions.
Deterministic	Questions in this category derive how deterministic the output for the function is required to be.
Cost Questions - 15	
Usefulness	Questions in this category elicit how useful automating the function would be.
Time	Questions in this category elicit how much time might be saved by automating the function
Criticality	Questions in this category request how flight or safety critical the function is. The premise is that these functions may cost more to automate.
Costs	Questions on cost range from # of lines of code to how long it might take to implement the function in software.
Efficiency/Task Mgt	To what degree would automating this function increase the efficiency of a human?
Mental Workload	To what degree would automating this function decrease a human's mental workload?
Boredom	Questions in this category elicit how repetitious is the function; the more repetitious, the more benefit automation might have

B. ADAPTATION OF FLOAAT FOR A UAS

Having described the general structure of FLOAAT and aspects of its derivation, this section describes how the tool was adapted for application to an unmanned aerial system. The first step was to determine a set of functions to be scored. The second step was to evaluate the rating scale for applicability, and adapt it where necessary. The third step included examining the NASA FLOAAT questions to ensure transferability to a UAS and to include questions that established architectural considerations that would enhance the resilience of the system. The application of these adaptations is detailed in Chapter IV as applied to a set of UAS functions.

1. Functional Architecture of an Unmanned Aerial System

Before delving into function selection to assess which level of autonomy is appropriate for that function, it is important to present a general functional architecture of a UAS. There does not exist a standard UAS functional architecture. However, the Office of the Secretary of Defense (OSD) has commissioned a cross-DOD project to build a common UAS control capability (UAS Control Segment (UCS) Working Group 2014). This thesis presents and references aspects of the OSD architecture to represent community level commonality and wide audience consumption.

An unmanned aerial vehicle (UAV) system consists of numerous subsystems and components that must seamlessly interact in order to meet its objectives. The payload is a complex system of systems in and of itself, often with numerous complex operating modes that generates considerable volumes of data in a short time. Communication links are required to transmit the current state of the vehicle as well as sensor data to mission or ground control so that operators can assess and evaluate the data. Mission control is provided through the a ground or surface based mission management capability that leverages communications to talk to the unmanned aircraft flight control and mission management computers. The mission control is the ground system that provides

the UAV operators with aircraft and environmental situational awareness, collaboration, and decision-making tools. It is used to support pre-flight planning, monitoring missions schedules, direct the payload system, visualization, and integrate sensing goals into the mission planning (Sullivan et al. 2004).

Figure 4 below depicts the top-level use cases performed by DOD UASs. The range is extensive, from strike to communications relay to moving cargo. “Perform Intelligence, Surveillance, and Reconnaissance (ISR) Mission Task” is circled in red, as that is the type of UAS employed as an example of function selection for this thesis.

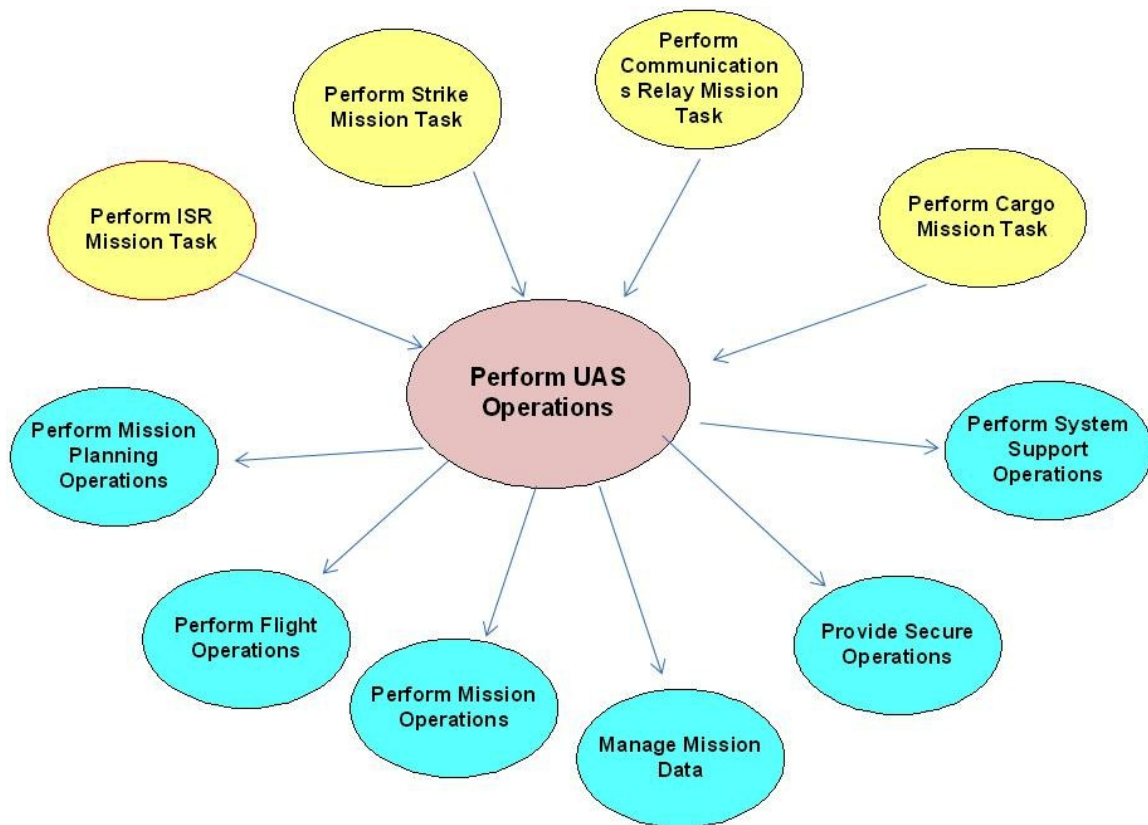


Figure 4. Contextual Use Cases for an Unmanned Control System
(from UCS Control Segment Working Group 2014)

The UAS Control Segment Architecture Description published by OSD contains extensive architectural artifacts on the overall UAS control system effort

and is pulling the community together to build out a control system that is open, modular, and that, eventually, can control different types of unmanned aerial vehicles. At present, most UAVs are slewed to the control station that was built when the air vehicle was built.

From the overall operational use case view, stakeholder analysis, and functional analysis, the following functional architecture for a UAS shown in Figures 5 and 6 is derived:

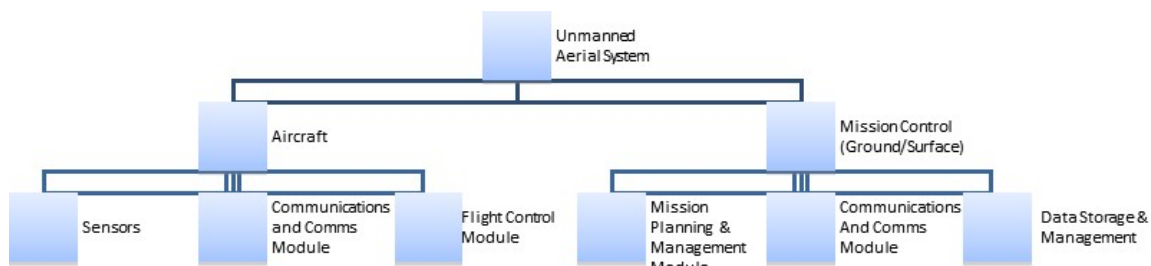


Figure 5. Functional Architecture of a UAS

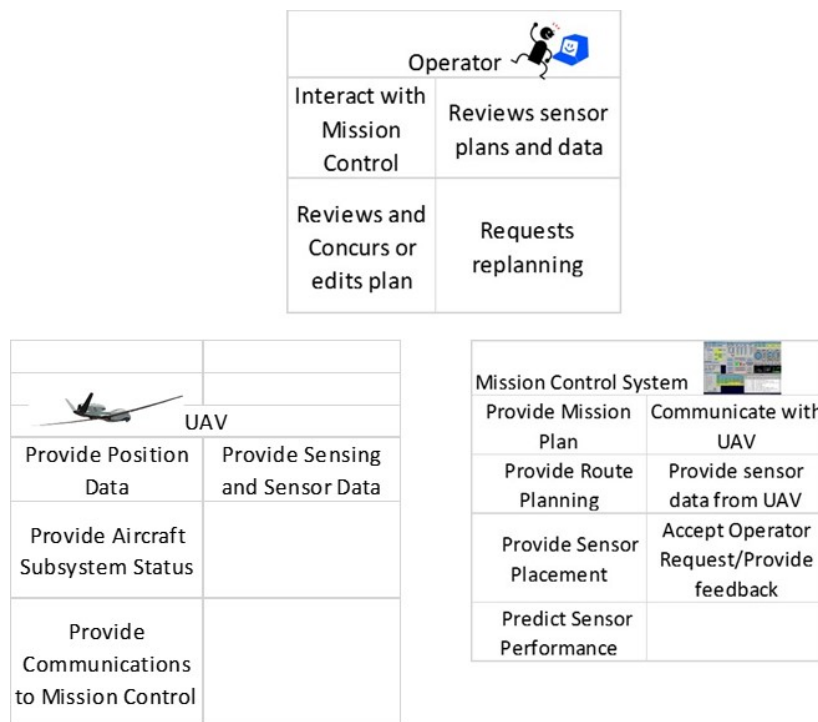


Figure 6. Functions Performed by UAS Segments

This process of decomposing top level functions and allocating them to a top level physical architecture helped frame the functions that were derived and employed in the analysis for this thesis. Table 3 depicts the functions derived, and provides and provides a short definition of each function. This process of system decomposition and definition is important, as it provides a frame of reference for the definitions of the functions. In Chapter IV, this same list will be presented along with the decision-making category, for example, Observe, Orient, Decide or Act, each function was assigned.

Table 3. A Set of UAS Mission Functions
(from UCS Working Group 2014)

Stage	Function Number	Function Name	Function Description
Pre-Planning		Mission Plan Provisioning	
		Mission Objectives	
	1	Objective Function Selection	Decide which objective to select in order to complete mission
	2	Nominal Take-Off, Cruise to Mission Area Flight Constraints	Determine mission constraints based on environmental and system limitations
	3	Flight Route Optimization Analysis (Against Sensing Objectives)	Determine, based on provided mission objectives, approach to route optimization
		Route Planning	
	4	Weather, Environment Data and Information	Research, retrieve environmental information related to the mission
	5	Vehicle/Flight Model Interpretation & Check	Retrieve appropriate model for air vehicle to enable evaluation of performance on recommended route
	6	Predict Take Off-On Mission Performance Margin	Provide potential performance measures of planned mission to assess the margin of performance
	7	Sensor System Evaluation	Provide evaluation of sensor system status/performance based on mission objectives and onboard sensor capabilities
	8	Flight Route Performance and Constraint Evaluation	Evaluate whether route is flyable based on known constraints and available mission avionics/fuel.
	9	Evaluate Mission Area Coverage	Evaluate how well mission area is covered by available sensors and flight capability
Take Off	10	Evaluate Flight Abort Coverage and Contingencies	Evaluate whether planned contingencies and aboard air bases are valid
	11	Route Optimization Decision	Determine/decide on optimal route for the mission.
	12	Mission/Flight Route Generation - Accept/Reject	Decide whether or not the planner recommended mission plan is acceptable or reject for further tweaking
		Take Off	
	13	Measure/Project Vehicle Conditions	Determine from onboard sensors status of air vehicle subsystems
	14	Predict On Mission Fuel Usage	Determine based on take-off factors what the fuel usage will be upon entering the mission area
	15	Current Flight Route Evaluation	Determine if planned route is still feasible/appropriate
	16	Alternative Flight Route Evaluation	Where desired/required, determine, review performance on an alternate route, as a means for comparison to baseline route
	17	Margin Calculation	
	18	Fuel Status Determination	Calculate/determine fuel state/availability
	19	Fuel Prediction	Predict if fuel is still adequate for mission accomplishment
	20	Take Off Abort Decision Execution	Make decision whether or not to aboard the mission
Execution		Flight Route Assessment	
	21	assess planned flight route/determine sensitivities	Upon mission area entry, reassess flight route plan and how immediate environment, mission situation affects the planned mission approach (route, sensor plan)
	22	resolve flight route conflicts	resolve any conflicts on route by recommending alternate, appropriate route/plan
	23	modify on mission objectives or rtn to base	determine if/how to modify mission objectives, or return to base
		ISR Mission Assessment	
	24	Review initial sensing objectives and constraints	Plot, assess sensing objectives and determine constraints
	25	Obtain sensor and sensing status	Retrieve sensor configuration, status, capabilities
	26	resolve sensing & route conflicts	Re-evaluate route based on sensor configuration and capabilities
	27	integrated plan assessment	Assess integrated sensor/route plan
	28	recommend modifications to mission objectives	recommend modifications to the plan based on mission performance, integrated plan information. Make updates to flight/mission plan.
		Landing/Recovery Opportunities Evaluator	
Landing	29	Landing Abort Information	Retrieve information to assist with decision on aborting the mission and landing safely
	30	Landing Site Validation	re-validate planned landing is still good.
		Contingencies	
	31	Pullout Calc & Assessment	Upon exiting mission area, determine what
	32	Landing/Recovery Update Monitoring (of systems)	Determine ability to land safely- provide margin (fuel, etc)
	33	Landing Location Recomputation	Recompute/revalidate landing location still valid
	34	Landing Action (or wave off, come back)	Determine pull out threshold and assess ability to recover

2. Adaptation of the Supervisory Control Reference Scale

Besides the fact that the tool was applied to a set of UAS functions vice spacecraft functions, a second modification applied was to the rating levels. NASA employed five levels. What ultimately was used for this research has ten. Two primary reasons influenced this decision: 1) Alignment with the scales of other supervisory control scales, to include the Cooper-Harper rating scale used within the world of aviation, and 2) to include increased detail for reference when determining the level of autonomy of a function.

As mentioned in Chapter II, the world of aviation is familiar with the Cooper-Harper scale. This scale was consulted as a factor for evaluating the levels as defined by NASA in FLOAAT. The Cooper-Harper rating scale is depicted below in Figure 7. Unlike the supervisory control scales presented in Chapter II, it focuses on aircraft handling qualities, not computer decision making. But, the frame of reference it provides is illustrative of what human-operators connote as a good aircraft system, which is worth consulting in informing a scale that could be applied to a UAS and the set of functions that comprise it.

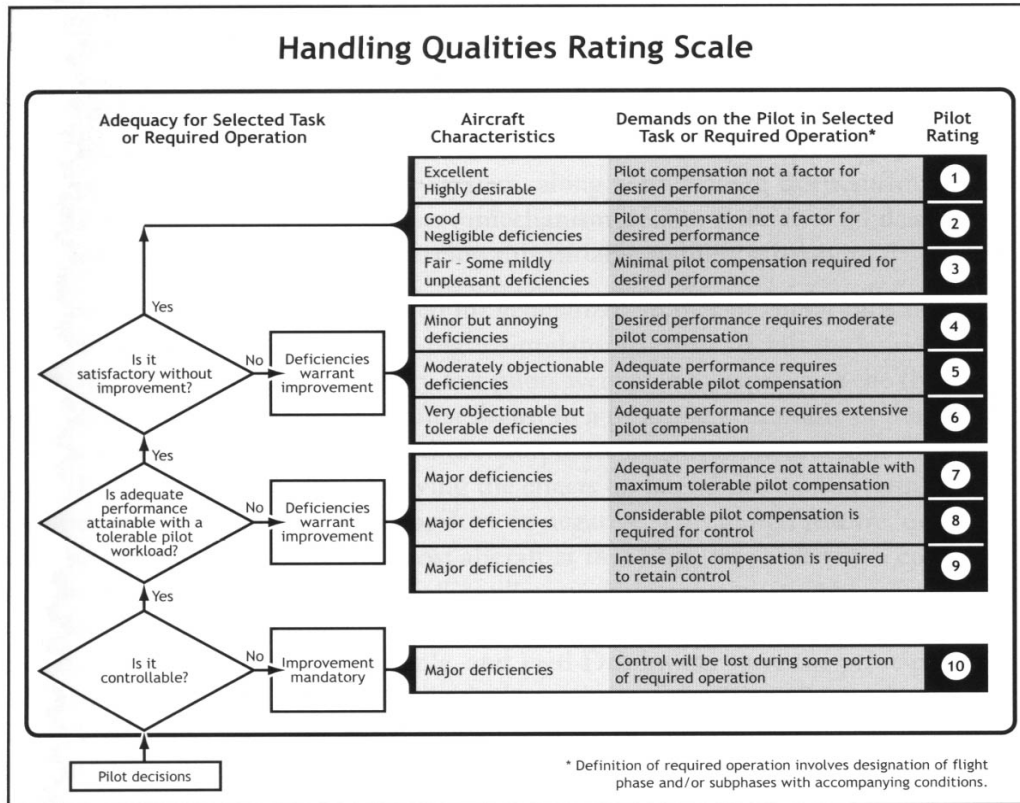


Figure 7. Cooper-Harper Handling Qualities Rating Scale
(from NASA History Web Site 2014b)

Though the Cooper-Harper scale could not be adopted wholesale, since it addresses a different purpose, one characteristic did apply, and that was the rating scale. The ratings numbered from one to ten, which just so happened to be the same as that devised by Sheridan and by Parasuraman (Sheridan and Johannesen 1976; Parasuraman 2000). It was decided to move to a 1–10 rating scale versus retaining the 5-level scale employed by NASA. This alignment enabled more direct refinement of what each level of autonomy could mean for a UAS, in within the OODA framework adopted by NASA. What the adaptation involved was repurposing the 5-level categorization for a spacecraft, a direct depiction shown in Table 4, to that of a 10-level scale for a UAS, shown in Table 5, both shown below.

Table 4. NASA's 5-pt Scale and Definition of the Levels of Autonomy
(from Proud and Hart 2005)

Level	Observe	Orient	Decide	Act
5	The data is monitored onboard without assistance from ground support	The calculations are performed onboard without assistance from ground support	The decision is made onboard without assistance from ground support	The task is executed onboard without assistance from ground support
4	The majority of the monitoring will be performed onboard with available assistance from ground support	The majority of the calculations will be performed onboard with available assistance from ground support	The decision will be performed onboard with available assistance from ground support	The task is executed onboard with available assistance from ground support
3	The data is monitored both onboard and on the ground.	The calculations are performed both onboard and on the ground.	The decision is made both onboard and on the ground and the final decision is negotiated between them.	The task is executed with both onboard and ground support
2	The majority of the monitoring will be performed by ground support with available assistance onboard	The majority of the calculations will be performed by ground support with available assistance onboard	The decision will be made by ground support with available assistance onboard	The task is executed by ground support with available assistance onboard
1	The data is monitored on the ground without assistance from onboard	The calculations are performed on the ground without assistance from onboard	The decision is made on the ground without assistance from onboard	The task is executed by ground support without assistance from onboard

If one reads even only one column or one row of definitions, one starts to appreciate the difference in definition between the one for a spacecraft and that for a UAS. The spacecraft definitions often allude to whether or not the ground control element should be informed or in control. The rating scale that numbers from one to ten also should be remembered as each UAS function will have a resulting score that falls with this range and aligns with how much control the UAS has versus the human-operator.

Table 5. Level of Autonomy Reference Rating Scale for a UAS
(from Proud and Hart, 2005)

Level	Observe	Orient	Decide	Act
10	UAS observes and monitors all systems and commands and acts autonomously, ignoring the human	UAS gathers data and information and interprets and integrates data and prepares to take action without involving the human-operator.	UAS performs analyses and ranks results for decision making, and does not display results to the human-operator.	UAS observes and monitors, analyzes, decides and acts autonomously, ignoring the human
9	UAS observes and monitors all systems and commands and acts autonomously, but informs the human after execution	UAS gathers data and information and interprets and integrates data and prepares to take action informing the human-operator but not waiting for consent. Does not display results	UAS performs analyses and ranks results for decision making, does not display results to the human-operator, but will upon query	UAS observes and monitors all systems and commands and acts autonomously, but informs the human after execution
8	The UAS gathers, filters, and prioritizes data; displays information only if asked	The UAS gathers data predicts, interprets, and integrates data into a result which is displayed to the human-operator only upon request	The UAS performs ranking tasks. The UAS performs final ranking, but does not display results to the human.	UAS executes automatically and does not allow any human interaction.
7	The UAS gathers, filters, and prioritizes data without displaying any information to mission control or the human. Status on command execution is provided, however.	The UAS analyzes, predicts, interprets, and integrates data into a result which is only displayed to the human if result fits programmed context (context dependant summaries).	The UAS performs ranking tasks. The UAS performs final ranking and displays a reduced set of ranked options without displaying "why" decisions were made to the human.	UAS executes automatically and only informs the human if required by context. It allows for override ability after execution. Human is shadow for contingencies.
6	The UAS gathers, filters, and prioritizes information displayed to the human.	The UAS overlays predictions with analysis and interprets the data. The human is shown all results.	The UAS performs ranking tasks and displays a reduced set of ranked options while displaying "why" decisions were made to the human.	UAS executes automatically, informs the human, and allows for override ability after execution. Human is shadow for
5	The UAS gathers information from the subsystems and environment, but it only displays non- prioritized, filtered information.	The UAS overlays predictions with analysis and interprets the data. The human shadows the interpretation for contingencies.	The UAS performs ranking tasks. All results, including "why" decisions were made, are displayed to the human.	UAS allows the human a context-dependant restricted time to veto before execution. Human shadows for contingencies.
4	The UAS and mission control is responsible for gathering the information for the human and for displaying all information, but it highlights the non-prioritized, relevant information for the user.	The UAS analyzes the data and makes predictions, though the human is responsible for interpretation of the data.	Both human and UAS perform ranking tasks, the results from the UAS are considered prime.	UAS allows the human a pre-programmed restricted time to veto before execution. Human shadows for contingencies.
3	The UAS is responsible for gathering and displaying unfiltered, unprioritized information for the human. The human still is the prime monitor for all information.	UAS is the prime source of analysis and predictions, with human shadow for contingencies. The human is responsible for interpretation of the data.	Both human and UAS perform ranking tasks, the results from the human are considered prime.	UAS executes decision after human approval. Human shadows for contingencies.
2	Human is the prime source for gathering and monitoring all data, with UAS shadow for emergencies.	Human is the prime source of analysis and predictions, with UAS shadow for contingencies. The human is responsible for interpretation of the data.	The human performs all ranking tasks, but the UAS can be used as a tool for assistance.	Human is the prime source of execution, with UAS/computer assistance for contingencies.
1	Human is the only source for gathering and monitoring (defined as filtering, prioritizing and understanding) all data.	Human is responsible for analyzing all data, making predictions, and interpretation of the data.	The UAS/mission control does not assist in or perform ranking tasks. Human must do it all.	Human alone can execute decision.

Each level for each decision-making category was redefined for a UAS leveraging the NASA scale, the Sheridan scale and the Parasuraman scale as an original reference. Ultimately, the scale is subjective; but, a careful read does show how each level is slightly different from the other as one moved from a low level of autonomy (levels 1–3) to a high level of autonomy (levels 8–10). In fact, the scale can be generally broken into three tiers: manual, autonomous with consent, and fully autonomous.

3. Inclusion of Architectural Attributes for a Resilient System

Before applying the FLOAAT set of 35 questions to UAS functions, each question was examined for applicability. The questions were also assessed as to whether they addressed resilience and, if not, if any appropriate wording could be added to enhance a designer’s thoughtfulness and consideration of architectural attributes that would lead to more resilient systems. Several enhancements were made. Table 6 presents the questions with modifications or enhancements shown in red font. Slight additions were made to certain categories in a way that would elicit consideration of resilience. The full set of questions with explanations for the questions are provided in Appendix A. Application of this questionnaire will be addressed in Chapter IV.

Table 6. Adaptation of FLOAAT to Consider Architectural Attributes of Resilient Systems (from Proud and Hart, 2005)

FLOAAT Questionnaire Adapted to include architecture attributes to improve resilience	
Ability	What is the expected ability of developers to correctly design the function for all possibilities within the design phase deadlines?
	What is the expected ability of programmers to correctly implement the design within the implementation deadlines?
Difficulty	What is the expected effort of developers to correctly design the function for all possibilities within the design phase deadlines?
	What is the expected effort of programmers to correctly implement the design within the implementation deadlines?
Robustness	What is the likelihood of an "outside-the-box" scenario occurring? How could the human-operator be a factor in functional redundancy to negotiate the "out of the box" scenario?
	How well will/can the function be designed to manage "outside-the-box" scenarios?
Experience	How autonomous (what level) has the function been shown to perform? Could the human-operator become too trustworthy and therefore become complacent? What functional redundancy would be required if so?
	Has the function been completed solely by a human during the flight phase itself?
Understandability	How understandable of a mental model of the function can a human create, including how the function works, what the output means, how to interact with the function?
	What is the level of human understanding required to accurately decide when an override is necessary?
	If an override is performed, what is the ability of a human to come up with a solution themselves?
Art vs Science	How much would a human have to infer what the computer "really meant" or what the computer will do in the future?
Familiarity	How familiar, friendly, and natural will the output feel to the user?
Correctness	What is the probability that the computer could come up with an answer that is "more accurate" than a human?
Training	How much training would be required for a human to perform this function instead of performing the function highly autonomously?
	How much training would be required for a human to interface with a tool using this function based on current understanding of the implementation of this function? How can this training on the interface improve human response to improve adaptability?
	How much verification would be required for this function to be trusted to perform fully autonomously?
Override	Is an override capability required (yes or no)?
Deterministic	How deterministic is the output from this function? Can decision making be distributed between the computer and human-operator to improve resilience?
2 LOA Cost/Benefit Limit	
Usefulness	How critical is this function to an overall Autonomous Mission and Flight Management system?
	How useful would automating this function be?
Time	How much time is available to perform function, considering flight phase, circumstances, possible contingencies, etc.?
Criticality	What is the criticality of this function to vehicle safety?
	What is the criticality of this function to crew safety?
Costs	How many lines of code are expected? low <= 1000 med-low <= 10,000 med <= 50,000 med-high <= 100,000 high >100,000
	How much time to design the function is expected?
	How much time to implement the software for this function is expected?
	What is the level of required verification and validation?
	What is the required skill level of software writers?
Efficiency/Task Mgt	To what degree would automating this function increase the efficiency of a human?
	To what degree would automating this function decrease a human's mental workload? Are there approaches to automating this function that would enhance operator flexibility?
Mental Workload	
Boredom	How repetitious is the function (level of frequency)?
	How mundane (does not utilize the skills of the operator) is the function?

THIS PAGE INTENTIONALLY LEFT BLANK

IV. APPLICATION OF THE FUNCTION-SPECIFIC LEVEL OF AUTONOMY AND AUTOMATION TOOL TO AN UNMANNED AIR SYSTEM

Having described the methodology of FLOAAT and the logic and analysis that underpins it, as well as adaptations that were made, this section presents how the tool was applied to a set of UAS functions and the resulting lessons learned. The application of the tool followed the same steps that NASA employed. The steps are as follows: 1) categorize each function according to the decision-making framework; 2) for each function, answer the 35-question questionnaire; 3) collect the resulting score and compare with the OODA framework for a second reference and check; 4) collect the scores for all the functions and evaluate them in a summary form.

A. CATEGORIZING UAS FUNCTIONS INTO THE OODA FRAMEWORK

Table 7 below depicts the list of 34 functions that were scored for the level of autonomy required. This figure is also provided in Appendix B as an embedded Microsoft Excel file in case additional detail is desired. The functions are listed by four main phases of a typical mission—pre-mission planning, take off, mission execution, then landing and include a short definition of the function in the right column. Pre-planning is divided into establishment of mission objectives and route planning. Mission Execution consists of Flight Route Assessment and ISR Mission Assessment. Examination of the list shows that some functions are as simple as obtaining status of avionics, to more complex functions that provide an optimization for the mission or a segment of the mission.

This list of functions is the same that were derived through systems decomposition and functional analysis shown in Chapter III (Table 3), except that in this case each is color-coded. The coloring of each function depicts how each was categorized according to the OODA framework. The color code is at the top of the table, with each of the four categories assigned a different color. The purpose for this categorization is so that the appropriate rating level and

interpretation is employed as reference (see Table 5 in Chapter III) when filling out the Level of Autonomy Questionnaire.

For example, the levels in the “Observe” column refer to gathering, monitoring, and filtering data; the levels in the “Orient” column refer to deriving a list of options through analysis, trend prediction, interpretation and integration; the levels in the “Decide” column refer to decision-making based on ranking available options; and the levels in the “Act” column refer to how autonomously the UAS takes action on a chosen option.

The intent is to help system designers to identify the level of autonomy for a function. When the levels for each category were redefined to ten rating levels, care was taken to provide more specific verbal description for each level.

Table 7. UAS Functions defined and categorized

Observe	Orient	Decide	Act
Stage	Function Number	Function Name	Function Description
Pre-Planning		Mission Plan Provisioning	
		Mission Objectives	
	1	Objective Function Selection	Decide which objective to select in order to complete mission
	2	Nominal Take-Off, Cruise to Mission Area Flight Constraints	Determine mission constraints based on environmental and system limitations
	3	Flight Route Optimization Analysis (Against Sensing Objectives)	Determine, based on provided mission objectives, approach to route optimization
		Route Planning	
	4	Weather, Environment Data and Information	Research, retrieve environmental information related to the mission
	5	Vehicle/Flight Model Interpretation & Check	Retrieve appropriate model for air vehicle to enable evaluation of performance on recommended route
	6	Predict Take Off-On Mission Performance Margin	Provide potential performance measures of planned mission to assess the margin of performance
	7	Sensor System Evaluation	Provide evaluation of sensor system status/performance based on mission objectives and onboard sensor capabilities
	8	Flight Route Performance and Constraint Evaluation	Evaluate whether route is flyable based on known constraints and available mission avionics/fuel.
	9	Evaluate Mission Area Coverage	Evaluate how well mission area is covered by available sensors and flight capability
	10	Evaluate Flight Abort Coverage and Contingencies	Evaluate whether planned contingencies and aboard air bases are valid
	11	Route Optimization Decision	Determine/decide on optimal route for the mission.
	12	Mission/Flight Route Generation - Accept/Reject	Decide whether or not the planner recommended mission plan is acceptable or reject for further tweaking
Take Off		Take Off	
	13	Measure/Project Vehicle Conditions	Determine from onboard sensors status of air vehicle subsystems
	14	Predict On Mission Fuel Usage	Determine based on take-off factors what the fuel usage will be upon entering the mission area
	15	Current Flight Route Evaluation	Determine if planned route is still feasible/appropriate
	16	Alternative Flight Route Evaluation	Where desired/required, determine, review performance on an alternate route, as a means for comparison to baseline route
	17	Margin Calculation	
	18	Fuel Status Determination	Calculate/determine fuel state/availability
	19	Fuel Prediction	Predict if fuel is still adequate for mission accomplishment
	20	Take Off Abort Decision Execution	Make decision whether or not to aboard the mission
Execution		Flight Route Assessment	
	21	assess planned flight route/determine sensitivities	Upon mission area entry, reassess flight route plan and how immediate environment, mission situation affects the planned mission approach (route, sensor plan)
	22	resolve flight route conflicts	resolve any conflicts on route by recommending alternate, appropriate route/plan
	23	modify on mission objectives or rtn to base	determine if/how to modify mission objectives, or return to base
		ISR Mission Assessment	
	24	Review initial sensing objectives and constraints	Plot, assess sensing objectives and determine constraints
	25	Obtain sensor and sensing status	Retrieve sensor configuration, status, capabilities
	26	resolve sensing & route conflicts	Re-evaluate route based on sensor configuration and capabilities
	27	integrated plan assessment	Assess integrated sensor/route plan
	28	recommend modifications to mission objectives	recommend modifications to the plan based on mission performance, integrated plan information. Make updates to flight/mission plan.
Landing		Landing/Recovery Opportunities Evaluator	
	29	Landing Abort Information	Retrieve information to assist with decision on aborting the mission and landing safely
	30	Landing Site Validation	re-validate planned landing is still good.
		Contingencies	
	31	Pullout Calc & Assessment	Upon exiting mission area, determine what
	32	Landing/Recovery Update Monitoring (of systems)	Determine ability to land safely- provide margin (fuel, etc)
	33	Landing Location Recomputation	Recompute/revalidate landing location still valid
	34	Landing Action (or wave off, come back)	Determine pull out threshold and assess ability to recover

B. ASSESSING UAS FUNCTIONS USING THE FLOAAT QUESTIONNAIRE

Having determined and categorized the list of functions that should be evaluated, the next step is to answer the set of 35 questions, for each and every function, that get at the heart of trust and cost/benefit. Full examples that show comments to the questions posed for Weather and Mission Objective Assessment are provided in Appendix C. The two functions differ in that one scored high for autonomous management of that function (weather data retrieval and assessment), and the other scored lower (mission optimization decision).

Answering the questionnaire took at least thirty minutes per function. The questions forced contemplation of the function, its design aspects, complexity, and, depending on the function, elicited a response that at times was informed by emotion as well as rationality. For example, the “Weather, Environmental Data and Information” function appears simple on the surface. Every mission requires weather data to plan before takeoff and requires near real time weather surrounding the aircraft during the mission to maintain safe flight. Table 8 below shows the question comments and notes for this function. In general, the Question Notes column was not altered from the NASA questionnaire, as it provided needed context and amplification of the question posed in the column to the left of the scores. However, comments were provided for nearly each response to explain the score provided.

Table 8. Questions from the FLOAAT Questionnaire and Comments to Those Questions for the Weather and Environmental Data Function (adapted from Proud and Hart, 2005)

	LoA Scale	10	9	8	7	6	5	4	3	2	1	Question notes	Comments/Notes to
Ability	What is the expected ability of developers to correctly design the function for all possibilities within the design phase deadlines?			1								Expected ability of designers to completely define the world of possibilities that this function will face, before the final deadline. Ability is defined as able to do the job, not the designer's ability level.	
	What is the expected ability of programmers to correctly implement the design within the implementation deadlines?				1							Expected ability of software writers to completely code the design that the developers handed them, regardless of the size of the world that was defined in the design phase, before the final deadline. Ability is defined as able to do the job, not the programmer's ability level.	weather data is available and easily fed into systems. Flight models can adjust to the weather elements but some information is not detailed enough to provide a true projection of what the weather situation will be.

Scoring for each function is straight forward. The number of marks is tallied for each rating level, and then averaged using equal weighting. For the Weather function, this resulted in a score of 5.63 (shown in Table 9 below) for Trust and 6.14 for Cost/Benefit. Table 9 is not directly from Proud and Hart's work, as it is composed of adapted input and applies to a UAS as opposed to a space craft. The FLOAAT example published and represented in Appendix B does not have a weather function.

Table 9. Weather Function (from Proud and Hart, 2005)

Function Name	Scale Type (Ob, Or, D, or A)										
	Observe										
	Weather, Environment Data and Information										
	Question --> Answer 1 in most applicable column										
1	LOA Trust Limit	High					Low				
	LoA Scale	10	9	8	7	6	5	4	3	2	1
		0	1	2	2	5	5	2	1	1	0
Weights		0	0.053	0.105	0.105	0.263	0.263	0.105	0.053	0.053	0
Absolute Scores		10	9	8	7	6	5	4	3	2	1
		Score					5.63				
2	LOA Cost/Benefit Limit										
		0	1	2	3	4	2	0	2	0	0
Weights		0	0.071	0.143	0.214	0.286	0.143	0	0.143	0	0
Absolute Scores		10	9	8	7	6	5	4	3	2	1
		Ave					6.14				

The Trust Score is lower than the Cost/Benefit Score. What this means is that there is a cost benefit to automating the function, but trust issues prevent a higher level of autonomy from being applied. The difference between the two scores is not great; both fall into the category that human-operator consent is still desired before auctioning this function. But, this research has shown that the details should be retained. The act of going through the thought process is what is significant. The comments and notes annotated along the way can serve as design artifacts for reference when the process of lower level design activities commence.

C. REVIEWING THE RESULTS IN SUMMARY

The answers to the 35 questions result in a composite score for Trust and a composite score for Cost that are shown in the Level of Autonomy (LoA) Trust Limit and LoA Cost Limit in the right most columns of Table 10. One step that is enforced in the overall process is the comparison of Trust to Cost/Benefit of automating that function. There is a clear difference between how much a human user trusts for a certain function to be handled by an autonomous system and how high the cost is to implement it. If the human-operator does not trust the

system, then it does not matter how intelligent or cost-efficient the system is designed to be. Similarly, even though a system would be highly trusted to work fully autonomously, there is no guarantee that this is the most cost-effective method of performing the function. A specific example is automated mission (flight route) validation for an unmanned air system. At present, a human operator performing the function is actually quicker and, therefore, more cost effective, than automating the function (Absil 2014). Given this premise, the highest level of autonomy is the minimum of how much a function scored in Trust or Cost.

Table 10. Level of Autonomy Scores for UAS Functions
(after Proud and Hard 2005)

Observe	Orient	Decide	Act	Design Phase: 2015-2018	
Stage	Function Number	Function Name	LOA Trust Limit	LOA C/B (Cost/Benefit) Limit	Min
Pre-Planning		Mission Plan Provisioning			
		Mission Objectives			
	1	Objective Function Selection	4.34	5.58	4.34
	2	Nominal Take-Off, Cruise to Mission Area Flight Constraints	5.80	4.90	4.90
	3	Flight Route Optimization Analysis (Against Sensing Objectives)	5.78	5.98	5.78
		Route Planning			
	4	Weather, Environment Data and Information	5.63	6.14	5.63
	5	Vehicle/Flight Model Interpretation & Check	5.05	5.50	5.05
	6	Predict Take Off-On Mission Performance Margin	6.05	6.50	6.05
	7	Sensor System Evaluation	7.12	7.56	7.12
	8	Flight Route Performance and Constraint Evaluation	5.57	5.71	5.57
	9	Evaluate Mission Area Coverage	5.10	5.90	5.10
Take Off	10	Evaluate Flight Abort Coverage and Contingencies	4.89	5.93	4.89
	11	Route Optimization Decision	4.79	5.79	4.79
	12	Mission/Flight Route Generation - Accept/Reject	4.33	5.21	4.33
		Take Off			
	13	Measure/Project Vehicle Conditions	6.23	5.98	5.98
	14	Predict On Mission Fuel Usage	6.12	5.85	5.85
	15	Current Flight Route Evaluation	5.34	5.58	5.34
	16	Alternative Flight Route Evaluation	5.38	5.85	5.38
	17	Margin Calculation	6.11	5.98	5.98
	18	Fuel Status Determination	7.62	7.88	7.62
	19	Fuel Prediction	6.11	6.02	6.02
	20	Take Off Abort Decision Execution	4.89	5.23	4.89
Mission Execution		Flight Route Assessment			
	21	assess planned flight route/determine sensitivities	5.23	5.11	5.11
	22	resolve flight route conflicts	4.42	4.97	4.42
	23	modify on mission objectives or rtn to base	4.14	4.87	4.14
		ISR Mission Assessment			0.00
	24	Review initial sensing objectives and constraints	5.43	5.55	5.43
	25	Obtain sensor and sensing status	6.90	6.53	6.53
	26	resolve sensing & route conflicts	5.85	5.47	5.47
	27	integrated plan assessment	5.12	5.42	5.12
	28	recommend modifications to mission objectives to VMS	5.34	5.12	5.12
Landing		Landing/Recovery Opportunities Evaluator			
	29	Landing Abort Information	5.99	6.12	5.99
	30	Landing Site Selection/Validation	4.77	5.23	4.77
		Contingencies			
	31	Pullout Calc & Assessment	5.28	5.55	5.28
	32	Landing/Recovery Update Monitoring (of systems)	5.53	6.53	5.53
	33	LandingLocation Recomputation	5.38	5.47	5.38
	34	Landing Action (or wave off, come back)	4.12	5.12	4.12

Scores range from the low 4s to just over 7 on a scale from one to ten. This range of scores generally equates to levels of automation where the system does heaving lifting but still requires consent or confirmation from the human-operator. Recall that in Chapter III it was noted that the scale depicts high,

medium, and low autonomy levels. These three tiers generally equate to manual, autonomous with consent, and fully autonomous control (high). Table 11 below depicts these tiers and the general descriptions of autonomous behavior exhibited by the UAS.

Table 11. Summary Levels of Supervisory Control for Adjustable Autonomy (from Sheridan 1992 and Parasuraman 2000)

High Autonomy	Level	Description of Autonomy
Executive Autonomous Operations (Sheridan, Parasuraman 6-10)	10	UAS observes and monitors all systems and commands and acts autonomously, informing the human operator after the fact, displaying information only if asked. ignoring the human
	9	
	8	
	7	
Consent Based Autonomy (Sheridan, Parasuraman 3-5)	6	The UAS gathers, filters, and prioritizes information displayed to the human, in time for human-operator to provide consent or to intervene.
	5	
	4	
Manual Control (Sheridan, Parasuraman 1-3) Low Autonomy	3	The UAS is responsible for gathering and displaying unfiltered, unprioritized information for the human. The human still is the prime monitor for all information, responsible for filtering, prioritizing, and assessing the data).
	2	
	1	

What the scores indicate is that the human-operator still needs to and should be ready to engage at the appropriate time during a mission. When one examines the scores along with the OODA groupings, an interesting pattern is revealed. Those functions that fall into the Decide and Act categories tend to score slightly lower in the autonomy level. What this implies is that a higher level of control should be provided to the human-operator for these functions, likely because these functions involve a level of decision-making where the human-operator needs to remain “in-the-loop.” This observation is not yet decisive; additional exploration and research is required to determine if the observation is statistically significant. But, it is one to list for future research.

D. SUMMARY AND LESSONS LEARNED

In summary, FLOAAT proved to be an effective tool to get at the heart of what level of autonomy is appropriate for any one function. The approach forced thoughtful consideration of different design, employment and cost aspects of making a function autonomous, which, in a manner, forced thorough requirements analysis for that function. Employment of FLOAAT showed that the process for determining the Level of Autonomy for any one function is iterative; a subject matter expert, in working through the questions and rating level definitions wrestles with the derived level resulting from the tool, and ostensibly would negotiate the intent and meaning of this level with a broader systems design team.

Another reason FLOAAT is beneficial is that using the tool prevents an operator or subject matter expert from gaming the system. What is meant by this is that one cannot look at a function and its definition and assign, subjectively, what the rating level should be. The level is derived from answering the set of 35 questions, which results in a composite score.

The fact that FLOAAT enables systems designers to understand the level of autonomy to design into a system so as to engender and retain trust, provides the foundation in which a human-operator is better positioned to operate a system effectively. Using the framework provided by Vaneman and Triantis on architectural attributes for resilient systems, it provides a mechanism to define a loose-coupling between the human-operator's control needs and a system's ability to execute autonomously. It provides a means to define and design in various levels of autonomy that can be changed and altered by the human-operator when required. In this way, it provides "procedural flexibility," thereby designing in a means to adapt the system's operations when needed. It positions that operator to leverage their knowledge, intuition and experience to be able to act or react to unanticipated events, thereby improving the resilience of a system.

Though the tool has proven useful in initial research, further investigation is required to truly validate its employment in the UAS domain. NASA has applied this to several programs. In doing so, they have developed approaches to validate the level of autonomy as suggested by FLOAAT (Proud and Hart 2005). They have a baseline of experience to draw from. This is not the case for unmanned aerial systems. If this tool were to be more widely adopted, there is more work to be done:

1. Determination of the composition of the team who should participate in the process of defining the level of autonomy by answering the questionnaire. How many and of what type of subject matter experts should be included?
2. The Level of Autonomy tool employed and adapted was originally designed to ascertain the division of labor between the computer and the human-operator (Proud and Hart 2005). Additional questions could be added to determine the division of labor between what should be on the aircraft and what functions should be executed in the mission control system.
3. The questions should be refined and tested against a larger cross-section of users or subject matter experts to ensure the question is clear and the intent is communicated.
4. Test cases should be developed in order to more quickly validate the scores and even prototype the output.

THIS PAGE INTENTIONALLY LEFT BLANK

V. FLOAAT APPLICABILITY ASSESSMENT AND FUTURE RESEARCH

Adjustable Autonomy improves a system's ability to adapt. It does so by providing a level of autonomy along a continuum that can be dialed up or down in terms of level of automation. This directly allows the system to *adapt* to unanticipated obstacles or unforeseen events by leveraging the higher order cognition of the human-operator. The challenge, at times has been, how much autonomy to design into a system's certain functions. This research contemplates the levels of autonomy that can be designed in for adjustment so as to improve the resilience of a system. Properly defined levels of autonomy directly addresses the trust issues of a human-operator, and optimizes the performance of the human and performance of the autonomous system. Regardless of how much autonomy is implemented in a system, some level of human-operator involvement will be required in interacting with that system (Glas and Kanda 2012). This is the case even if the only task the human-operator has is to monitor the system, just in case he/she needs to terminate a certain operation. That the level of autonomy can be adjusted is important. For the application to the unmanned aerial system on an ISR mission, this could be true when the human-operator gets busy handling or acknowledging sensing aspects of the mission and needs to leave flight handling completely to the system to take care of. Perhaps, the mission is long and boring. In this case, most functions can be set at fully autonomous, with the feature that the system alerts the human if it senses the mission, the environment, its operations require attention. Given this premise, the human-system roles and interface should be defined and designed so that the human-operator is able to graduate to more autonomous operation as the trust level increases.

A. EVALUATION OF ADAPTING FLOAAT TO A UAS

Having adapted and then applied NASA's FLOAAT to a set of UAS functions, the following points can be made:

1. The FLOAAT Questionnaire

As had been mentioned, the wording of the questions should be further tested against a broader group of relevant subject matter experts who have different roles and, therefore, different views. An attempt was made to get additional personnel to answer the questions, but the time this would require was too extensive. To properly fill out the questionnaire for a function took at least 30 minutes; 34 questions would require at least 15 hours per person. Due to resource and time limitations, this bank of questions was answered by only one person. This person, the author, has nearly ten years of experience working on fielded and unfielded unmanned aerial systems; still, the answers are from only one person, from one perspective.

Due to the subjective nature of determining the appropriate level of autonomy in designing a system, personal biases and technical experiences could affect the results. The questions are interpreted differently from person to person. The wording of subjective questions, corresponding notes, and examples in this tool is critical, especially if the question refers to a length of time or milestone (which must be explicitly stated).

Besides testing the wording of the questions further, the size and composition of the group of respondents needs to be determined. How many need to answer to ensure relevance? What backgrounds should be sought out? Obviously, the more that answers the questions, the better. But, such an approach could become costly. At the very least, respondents need to be from groups who have true input and impact on the system. NASA had sought out a low number, trusting subject matter expertise and engineering judgment, but they did seek out various backgrounds—flight controllers, trainer, ground control operators, systems engineers (Proud and Hart 2005).

Finally, what should not be lost when examining scores provided to any one function are the comments that are provided along with the numerical score. The questionnaire provided an additional column for users to provide comments

and amplification for their scores. Oftentimes the comments provided the underpinning for the answer on how their score aligned with the OODA framework, or how they may have interpreted the question when applied to the function at hand. Other times comments suggested design approaches that addressed architectural attributes, like how to design in adaptability to a certain function.

2. Employment of OODA Categories

Contemplating the process of using the FLOAAT process in defining a level of autonomy brings up the question of whether or not bucketing each function in the OODA decision categories was useful or not. On face value, the process seemed to be an extra step that took time and the definitions for each category does not seem to be that different. However, when contemplating a single function, the framework does become useful. It forces a thought process that helps further identify the level of autonomy the person answering the question is comfortable with, or is certain of, and provides a means of referencing that level to a scale. NASA had created these categories and differentiated the rating levels; a level 5 in Observe is not the same as a level 5 in Orient, Decide or Act. Their work has created a foundation that has solidified these definitions. For this tool to be truly useful, the same needs to be done for the UAS community. However, even if this step were not taken, in the end, the OODA framework proved useful. Ultimately, it served as a validation check on whether or not the resulting score for any one function made sense.

3. Employment of the 10-Level Rating Scale

As had been pointed out in preceding Chapters, despite the existence of several levels of 10-level rating scales, autonomy levels can generally be bucketed into three tiers: fully autonomous, autonomous with consent, and manual. Therefore, there arises the question on whether or not a 10-level scale is necessary or that a 3-level rating scale is sufficient? Would not a simpler scale reduce complexity and be easier to employ? Based on the experience of going

through the process of using the 10-level rating scale, the answer is that for designers and for subject matter experts answering the questions, the granularity is helpful. Though the output from this research has not been employed in actually implementing a level of autonomy using this system, which would validate some of the statements made in this conclusion, the detailed levels have already assisted with requirements analysis and definition. Potentially, the conclusion should be that this tool should be employed by systems engineers to obtain the Voice of the Customer in order to further refine system level requirements in order to develop and interface that is suitable and trusted by the human-operator.

4. Incorporation of Architecting for Resilience

The questionnaire, in many respects, helps a systems engineer/designer wrestle with the intent of a systems level requirement. In this way, it facilitates requirements analysis. There is an opportunity to insert questions or modify questions to ensure certain architecting principles are included, such as those which would ensure consideration of resilience principles—loose coupling, means to enhance robustness, functional and physical redundancy, etc. This was done in a modest fashion in this research and could be extended more broadly. Some of the questions inherently got at systems robustness.

B. ADDITIONAL CONSIDERATIONS AND FUTURE RESEARCH

There is a key aspect that should be kept in mind and explored as the military and operators get more accustomed to unmanned systems and increased levels of autonomy. Experience shows that human operators have a tendency to become reliant on the automated/autonomous systems. This can be particularly challenging during system degradation, sometimes without the full awareness of the human-operator. Overreliance on automation is frequently suspected as a factor in or indeed the cause of aviation incidents and accidents. Overreliance and loss of operator proficiency can result in human-operators

becoming reluctant to assume manual control even when the autonomous system is not operating correctly. (National Research Council 2014).

Autonomous systems have been changing the way the military does business, and, with recent investment by the DOD and the commercial world, is on the threshold of exerting deep changes in military operations. These systems can and will be able to be operated without direct human control for extended periods of time and over long distances. This is beneficial and will open the field for more applications while reducing costs; but, it should be done with the human-operator, and his/her strengths and weaknesses, in mind. Or else, the systems may not be adopted, or, even worse, the systems may not be safe. As such, the following are a few suggested areas of further research:

1. Human-Operator Collaboration

- Determine how the roles of human-operations and the autonomous systems, as well as the human-system interface, should evolve to enhance more efficient yet safe operations.
- Further understanding of human psychology in the operation of autonomous systems.
- Interfaces, be they visual, aural, focused on assistance or alerting to problems that improve human performance.
- Approaches to adjust to different skill and cognition levels in human-operators, with an eye toward safety.

2. Verification, Validation, and Certification

- Develop standards and processes for the verification, validation, and certification of autonomous systems, and determine their implications for architecture and design.

3. Autonomy Architecture

- Explore and define the landscape of autonomous systems architecture to further the ability to adapt and verify and validate the system.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX A. LEVEL OF AUTONOMY QUESTIONNAIRE

Table 12. Level of Autonomy Questionnaire (from Proud and Hart, 2005)

Function Name	Scale Type (Ob, Or, D, or A)		
	Observe		
	Function X		
	Question --> Answer 1 in most applicable column		
1	LOA Trust Questions		
	LoA Scale	Question notes	Comments and Notes to Score Provided
Ability	What is the expected ability of developers to correctly design the function for all possibilities within the design phase deadlines?	Expected ability of designers to completely define the world of possibilities that this function will face, before the final deadline. Ability is defined as able to do the job, not the designer's ability level.	
	What is the expected ability of programmers to correctly implement the design within the implementation deadlines?	Expected ability of software writers to completely code the design that the developers handed them, regardless of the size of the world that was defined in the design phase, before the final deadline. Ability is defined as able to do the job, not the programmer's ability level.	

Function Name	Scale Type (Ob, Or, D, or A)		
Difficulty	What is the expected effort of developers to correctly design the function for all possibilities within the design phase deadlines?	This is the same as the above questions, but the focus is not on "how good will the design be?" but on "how hard will it be to design?"	
	What is the expected effort of programmers to correctly implement the design within the implementation deadlines?	Focus is "how hard" - the coding of the selection function is straightforward; the development of math models, or characterization of what is needed is what is hard.	
Robustness	What is the likelihood of an "outside-the-box" scenario occurring?	Weather -always unpredictable	
	How well will/can the function be designed to manage "outside-the-box" scenarios?	Should be able to design the model; this question is not truly applicable	
Experience	How autonomous (what level) has the function been shown to perform?	Commercial and military systems have this function more or less automated - in mission planners/ground systems	
	Has the function been completed solely by a human during the flight phase itself?	This is a pre-mission function; humans may have acted to change the objective, and try and optimize, but likely was based on experience and gut feel	
Understandability	How understandable of a mental model of the function can a human create, including how the	Are the concepts, themselves, involved with this function complicated? Yes!	

Function Name	Scale Type (Ob, Or, D, or A)		
	function works, what the output means, how to interact with the function?		
	What is the level of human understanding required to accurately decide when an override is necessary?	What level of understanding would a human need to have in order to determine if the output from this function is out of family? It would be high. Humans tend to compensate via experience/heuristics	
	If an override is performed, what is the ability of a human to come up with a solution themselves?	Human will come up with solution, but may not be mathematically optimal	
Art vs Science	How much would a human have to infer what the computer "really meant" or what the computer will do in the future?	This is truly an Art vs. Science question. If performing this function is an art form of human fudge factors, and post-processing mental tweaks, then it should be hard to automate. Though if the function is purely scientific, with a definite answer that needs little human interaction to change it to be the "correct" answer, then the function should be easier to automate.	
Familiarity	How familiar, friendly, and natural will the output feel to the user?	Depends on how the designer/coder develops the presentation; perhaps potential answers could be provided visually for each objective function there is so the human	

Function Name	Scale Type (Ob, Or, D, or A)		
		can leverage experience to make a decision.	
Correctness	What is the probability that the computer could come up with an answer that is "more accurate" than a human?	Both a human and a computer can come up with an answer that is "right". A human may be able take the big picture and incorporate it into coming up with the better answer. A computer may be able to run optimization algorithms to come up with a better answer. "Better"- can be subjective	
Training	How much training would be required for a human to perform this function instead of performing the function highly autonomously?	Training would revolve around what the objective functions provided; what they mean to the mission	
	How much training would be required for a human to interface with a tool using this function based on current understanding of the implementation of this function?	Can a human do this task with some help from the computer?	
	How much verification would be required for this function to be trusted to perform fully autonomously?	how many cases and examples would have to be proven to work correctly for the function to be trusted to work fully	

Function Name	Scale Type (Ob, Or, D, or A)		
Override	Is an override capability required (yes or no)?	There may need to be with novice users who do not understand selection of a math model; This will limit the autonomy scale to allow override if yes is chosen	
Deterministic	How deterministic is the output from this function?	Flight constraints etc. tend to be fairly specific once flight model characterized	
Weights			
Absolute Scores			
2	LOA Cost/Benefit Questions		
Usefulness	How critical is this function to an overall Autonomous Mission and Flight Management system?	While the function itself might not be that critical, other functions might require this function to be done highly autonomously in order to work highly autonomously themselves.	
	How useful would automating this function be?	Gut feeling. This function would be useful for the computer to do instead of the human.	
Time	How much time is available to perform function, considering flight phase, circumstances, possible contingencies, etc.?	Each flight phase has a different scale of time. On Take Off and Landing phase, many decisions may be required in milliseconds. This is faster than a human could possibly provide an answer, trending	

Function Name	Scale Type (Ob, Or, D, or A)		
		towards more autonomy.	
Criticality	What is the criticality of this function to vehicle safety?	function would impact how vehicle would fly	
	What is the criticality of this function to crew safety?	System is unmanned; mission control needs to be protected, but there is no harm to crew as there is not a crew in the vehicle	
Costs	How many lines of code are expected? low <= 1000 med-low <= 10,000 med <= 50,000 med-high <= 100,000 high >100,000	Arbitrary scale based on a few conversations with Prakash Sarathy. Somewhat arbitrary assessment	
	** How much time to design the function is expected?	Question of man-hours. The deadline for completion is set, and this question asks will this function be done in time.	
	How much time to implement the software for this function is expected?	Same as previous question. Focused on writing the code, not the design phase.	
	What is the level of required verification and validation?	This is the software V&V question. How many runs will be needed to prove that the algorithms work.	
	** What is the required skill level of software writers?	How hard will the function be to program?	
Efficiency/Task Mgt	To what degree would automating this function increase the	Would this increase the efficiency of whoever interfaces with this	

Function Name	Scale Type (Ob, Or, D, or A)		
	efficiency of a human?	function, be it human or other function?	
Mental Workload	To what degree would automating this function decrease a human's mental workload?	Would the human still have to worry about this function? Could this be automated well-enough that the human does not have to think about it anymore.	
Boredom	How repetitious is the function (level of frequency)?	Answer based on the flight phase, and the number of cycles a second the function would be performed.	
	How mundane (does not utilize the skills of the operator) is the function?	Depends on the operator to some extent. But, if the task bores the human that is forced to perform it, the tendency is for errors to increase. Thus, mundane tasks should be automated.	

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX B: EXAMPLE OF NASA'S LEVEL OF AUTONOMY SCORES FOR A SPACECRAFT

Table 13. NASA's FLOAAT Scores for a Spacecraft Re-planning tool
(from Proud and Hart 2005)

Function Number	Function Name	LOA Trust Limit	LOA C/B (Cost/Benefit) Limit	Function type observe=1 orient =2 decide =3 act=4	Minimum
1 - Prelaunch	Ground Updates	4.00	5.63	1	4.00
2 - Prelaunch	Vehicle System Monitoring	4.00	5.38	1	4.00
3 - Prelaunch	Non-Trajectory-Related Flight Operations Monitoring	4.00	5.50	1	4.00
4 - Prelaunch	LW/LT Function Objective Determination	4.71	5.25	3	4.71
5 - Prelaunch	Selection of Nominal Insertion Target Altitude	5.22	5.75	3	5.22
6 - Prelaunch	Nominal Insertion Propellant Requirement	5.32	5.00	2	5.00
7 - Prelaunch	Monitor Launch Window expansion amount	4.00	5.50	1	4.00
8 - Prelaunch	Degraded Insertion Target Altitude - Launch Window Expansion	4.60	4.75	3	4.60
9 - Prelaunch	Launch Window Expansion Propellant Requirement	5.22	5.00	2	5.00
10 - Prelaunch	Predict Ascent Performance Margin	5.53	5.88	2	5.53
11 - Prelaunch	Degraded Insertion Target Altitude - Ascent Performance Loss	4.81	4.63	3	4.63
12 - Prelaunch	Ascent Performance Loss Propellant Requirement	5.22	5.00	2	5.00
13 - Prelaunch	Determine Phasing Windows	5.74	5.25	2	5.25
14 - Prelaunch	Determine Intersection Point or Point of Closest Approach	6.00	5.25	2	5.25
15 - Prelaunch	Determine Zero Wedge Angle Time (Analytic In-Plane Time)	6.00	6.00	2	6.00
16 - Prelaunch	Determine In-Plane Time	6.00	5.38	2	5.38
17 - Prelaunch	Determine Planar Window	6.00	5.75	2	5.75
18 - Prelaunch	Determine Planar/Phase Window Overlap	5.94	6.13	2	5.94
19 - Prelaunch	Evaluate Candidate LW/LTs Against Constraints	4.60	4.75	2	4.60
20 - Prelaunch	Rank Available Launch Targets	4.71	4.75	3	4.71
21 - Prelaunch	Store Available Launch Targets	6.35	7.13	4	6.35
22 - Prelaunch	Optimum Launch Target	5.32	6.00	4	5.32
23 - Prelaunch	Objective Function Selection	4.50	4.75	3	4.50
24 - Prelaunch	Nominal Ascent Flight Constraints	4.91	6.00	2	4.91
25 - Prelaunch	Planetary Environment Data	4.00	6.38	1	4.00
26 - Prelaunch	Vehicle Data	4.00	6.38	1	4.00
27 - Prelaunch	Trajectory Data	4.00	6.25	1	4.00
28 - Prelaunch	Staging Constraint Parameters	4.00	6.38	1	4.00
29 - Prelaunch	Mathematical Model of Objective Function	5.01	5.50	2	5.01
30 - Prelaunch	Selection of New Mathematical Model of Objective Function	5.12	5.50	3	5.12
31 - Prelaunch	Nonlinear Optimizer Tuning Properties Selection	5.74	6.25	3	5.74
32 - Prelaunch	Initial Trajectory Generation	4.91	5.00	4	4.91
33 - Prelaunch	Trajectory Optimization Analysis	4.81	5.88	2	4.81
34 - Prelaunch	Trajectory Optimization Decision	5.32	6.00	3	5.32
35 - Prelaunch	Feasibility and Convergence Check	5.22	6.13	2	5.22
36 - Prelaunch	Trajectory Performance and Constraint Evaluation	5.53	6.50	2	5.53
37 - Prelaunch	Optimized Trajectory	5.84	6.13	4	5.84

THIS PAGE INTENTIONALLY LEFT BLANK

BIBLIOGRAPHY

- Amalberti, Rene R. 1999. "Automation in Aviation: A Human Factors Perspective." In *Handbook of Aviation Human Factors*, edited by D. J. Garland, J. A. Wise, and V. D. Hopkin, 173–192. Mahwah, NJ: Erlbaum.
- Amalberti, Rene, and François Deblon. 1992. "Cognitive Modeling of Fighter Aircraft Process Control: A Step Towards an Intelligent On-Board Assistance System." *International Journal of Man-Machine Studies* 36, no. 5: 639–671.
- Bernard, Douglas E., Gregory A. Dorais, Charles Fry, Edward B. Gamble Jr., Bob Kanefsky, James Kurien, William Millar, and Nicola Muscettola. 1998. "Design of the Remote Agent Experiment for Spacecraft Autonomy." In *IEEE Aerospace Conference: Vol. 2*. Snowmass, CO: IEEE Aerospace Conference.
- Billings, Charles E. 1997. *Aviation Automation: The Search for a Human-Centered Approach*. Mahwah, NJ: Erlbaum.
- Billings, Charles E., and David D. Woods. 1994. "Concerns about Adaptive Automation in Aviation Systems." In *Human Performance in Automated Systems: Current Research and Trends*, edited by Mustapha Mouloua and R. Parasuraman, 24–29. Hillsdale, NJ: Erlbaum.
- Bonasso, R. P., D. Kortenkamp, and T. Whitney. 1997. "Using a Robot Control Architecture to Automate Space Station Shuttle Operations." Paper presented at the 14th National Conference on Artificial Intelligence and Ninth Conference on Innovative Applications of Artificial Intelligence (IAAI), Cambridge, Massachusetts, July 28-31.
- Bradshaw, Jeff M., Paul J. Feltovitch, Hyuckchul Jung, Shriniwas Kulkarni, William Taysom, and Andrzej Uszok 2004. "Dimensions of Adjustable Autonomy and Mixed-Initiatives Interaction." In *Agents and Computational Autonomy*, edited by M. Nickles, M. Rovatsos, and G. Weiss, 17–39. Berlin, Germany: Springer-Verlag.
- Boyd, John. R. 1996 "The Essence of Winning and Losing," Excerpts in presentation format found on the web at http://www.defense-and-society.org/fcs/ppt/boyds_ooda_loop.ppt.OODA.
- Christman, John. 2009. *Autonomy in Moral and Political Philosophy*. Stanford Encyclopedia of Philosophy. Last modified August 11. <http://plato.stanford.edu/entries/autonomy-moral>.

- Clough, Bruce. 2000. "Relating Autonomy to a Task - Can It Be Done?" Paper presented at the proceedings of the American Institute of Aeronautics and Astronautics (AIAA) 1st Technical Conference and Workshop on Unmanned Aerospace Vehicles, Portsmouth, Virginia, May 20–23, 2002.
- Commander, Naval Air Systems Command. 2008. *Naval Air Training and Operating Procedures Standardization (NATOPS) Flight Manual Navy Model F/A-18E/F 165533 and Up Aircraft (A1-F18EA-NFM-000)*. Patuxant River, MD: Naval Air Systems Command.
- Defense Science Board (DSB). 2012. *Task Force Report: The Role of Autonomy in DoD Systems*. Washington, DC: Office of the Under Secretary of Defense for Acquisition, Technology and Logistics, July 2012. <http://www.acq.osd.mil/dsb/reports/AutonomyReport.pdf>.
- Dorais, Gregory. A., R. Peter Bonasso, David Kortenkamp, Barney Pell, and Debra Schreckenghost. 1999. "Adjustable Autonomy for Human-Centered Autonomous Systems on Mars." Paper presented at the proceedings of the 6th International Joint Conference on Artificial Intelligence (IJCAI), Workshop on Adjustable Autonomy Systems.
- Dorais, Gregory. A., and David Kortenkamp. 2008. "Designing Human-Centered Centered Autonomous Agents." PowerPoint presentation. N.p.: National Aeronautics and Space Administration. <http://ti.arc.nasa.gov/m/pub-archive/189h/0189%20%28Dorais%29.pdf>.
- Duda, Richard. O and Edward. H. Shortliffe. 1983. "Expert Systems Research." *Science* 220, no. 4594: 261–268.
- Fiksel, Joseph. 2003. "Designing Resilient, Sustainable Systems." *Environmental Science and Technology* 37, no. 23: 5330–5339.
- Frost, Chad. 2010. "Challenges and Opportunities for Autonomous Systems in Space." Speech at the National Academy of Engineering's U.S. Frontiers of Engineering Symposium, Armonk, New York, September 23–24, 2010.
- Glas, Dylan F., Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. 2012. "Teleoperation of Multiple Social Robots." *IEEE Transactions on Systems, Man and Cybernetics, Part A*: 42, no. 3: 530–544.
- Goodrich, Michael A., Dan R. Olsen, Jacob W. Crandall, and Thomas J. Palmer. 2001. "Experiments in Adjustable Autonomy." Paper presented at the proceedings of the 17th International Joint Conference on Artificial Intelligence Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents, Seattle, Washington, August 4–10.

- Hart, Jeremy, Ryan Proud, and Joshua Hardy. 2005. *Autonomy and Automation Pilot Study: Rendezvous & Docking Final Report* (Technical Report AC1L1-17, AFM-FLOAAT0004). Houston, TX: NASA, Johnson Space Center.
- Hart, Jeremy, and John Valasek. 2007. Methodology for Prototyping Increased Levels of Automation for Spacecraft Rendezvous Functions. Lecture at Texas A&M University, May 9.
- Holling, Crawford S. 2009. "Engineering Resilience versus Ecological Resilience." In *Foundations of Ecological Resilience*, edited by Lance H. Gunderson, Craig R. Allen, and Crawford S. Holling, 51–65. Washington, DC: Island.
- Hollnagel, Erik, and David D. Woods. 2006. "Epilogue: Resilience Engineering Precepts." In *Resilience Engineering: Concepts and Precepts*, edited by Erik Hollnagel, David D. Woods, and Nancy Leveson, 347–357. Burlington, VT: Ashgate.
- Isby, David C. *Jane's How to Fly and Fight in the F/A-18 Hornet*. London: HarperCollins, 1997.
- Inagaki, Toshiyuki. 2003. "Adaptive Automation: Sharing and Trading of Control." In *Handbook of Cognitive Task Design*, edited by Erik Hollnagel, 147–169. Mahwah, NJ: Erlbaum.
- Jackson, Scott. 2010. *Architecting Resilient Systems: Accident Avoidance and Survival and Recovery from Disruptions*, edited by Andrew P. Sage. Hoboken, NJ: Wiley.
- Jackson, Scott, and Timothy L. J. Ferris. 2012. "Resilience Principles for Engineered Systems." *Systems Engineering* 16, no. 2: 152–164.
- Lenat, Douglas B., and Ramanathan V. Guha. 1989. *Building Large Knowledge-Based Systems*. Menlo Park, CA: Addison-Wesley.
- Lintern, Gavan, and Thomas Hughes. 2008. *Development of a Supervisory Control Rating Scale* (AFRL-RH-WP-TR-2008-0054). Wright-Patterson AFB, OH: Air Force Research Laboratory.
- Madni, Azad. M. 2007 "Designing for Resilience." ISTI Lecture Notes on Advanced Topics in Systems Engineering.
- Madni, Azad. M., and Scott Jackson. 2009. "Towards a Conceptual Framework for Resilience Engineering." *IEEE Systems Journal* 3, no. 2: 181–191.

- Malin, John, Carroll Thronesbery, and Debra Schreckenghost. 1996. Progress in Human-Centered Automation: Communicating Situation Information". Paper presented at the Conference on Life Sciences and Space Medicine, American Institute of Aeronautics and Astronautics (AIAA), Houston, TX, March 5-7, 1996.
- Martin, Cheryl, Debra Schreckenghost, Peter Bonasso, David Kortenkamp, Thomas Milam, and Carroll Thronesbery. 2003. "Aiding Collaboration Among Humans and Complex Software Agents." In *Association for the Advancement of Artificial Intelligence (AAAI) Spring Symposium on Human Interaction with Autonomous Systems in Complex Environments* (Technical Report SS-03-04). Stanford, CA: AAAI.
- Martin-Breen, Patrick, and J. Marty Anderies. 2011. *Resilience: A Literature Review*. New York: Rockefeller Foundation.
- National Research Council. 2014. *Autonomy Research for Civil Aviation: Toward a New Era of Flight*. Washington, DC: National Academies. http://www.nap.edu/openbook.php?record_id=18815.
- NASA History Web Site. 2014a. "SP-3300 Flight Research at Ames, 1940–1997." Last modified August 3, 2014. <http://history.nasa.gov/SP-3300/ch4.htm>.
- NASA History Web Site. 2014b. "Handling Qualities Rating Scale." Image. Last modified August 3, 2014. <http://history.nasa.gov/SP-3300/fig66.htm>.
- Parasuraman, Raja. 2000. "Designing Automation for Human Use: Empirical Studies and Quantitative Models." *Ergonomics* 43, no. 7: 931–951.
- Parasuraman, Raja, Toufik Bahri, John E. Deaton, Jeffrey G. Morrison, and Michael Barnes. 1992. *Theory and Design of Adaptive Automation in Aviation Systems*. Warminster, PA: Naval Air Warfare Center.
- Parasuraman, Raja, Galster, Scott, Squire, Peter, Furukawa, Hiroshi and Miller, Chris. 2005. "A Flexible Delegation-Type Interface Enhances System Performance in Human Supervision of Multiple Robots: Empirical Studies with RoboFlag". *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 35, no. 4: 481–493.
- Parasuraman, Raja, and Peter Hancock. 2001. "Adaptive Control of Mental Workload." In *Stress, Workload, and Fatigue*, edited by P. A. Hancock and P. A. Desmond, 305–320. Mahwah, NJ: Erlbaum.

- Parasuraman, Raja, and Christopher A. Miller. 2006. "Delegation Interfaces for Human Supervision of Multiple Unmanned Vehicles: Theory, Experiments, and Practical Applications." In *Human Factors of Remotely Piloted Vehicles*, vol. 7 of *Advances in Human Performance and Cognitive Engineering Research*, edited by Nancy J. Cooke, Heather L. Pringle, Harry K. Pedersen, and Olena Connor, 251–266. Amsterdam: Elsevier JAI.
- Parasuraman, Raja, and Mustapha Mouloua. 1996. *Automation and Human Performance: Theory and Application*. Mahwah, NJ: Erlbaum.
- Parasuraman, Raja, Mustapha Mouloua, and Brian Hilburn. "Adaptive aiding and adaptive task allocation enhance human-machine interaction." *Automation technology and human performance: Current research and trends* (1999): 119-123.
- Parasuraman, Raja, Mustapha Mouloua, and Robert Molloy. 1996. "Effects of Adaptive Task Allocation on Monitoring of Automated Systems." *Human Factors* 38, no. 4: 665–679.
- Parasuraman, Raja, and Victor Riley. 1997. "Humans and Automation: Use, Misuse, Disuse, Abuse." *Human Factors* 39, no. 2: 230–253.
- Parasuraman, Raja, Thomas B. Sheridan, and Christopher D. Wickens. 2000. "A Model for Types and Levels of Human Interaction with Automation." *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions* 30, no. 3: 286–297.
- Proud, Ryan W., and Jeremy J. Hart. 2005. "FLOAAT, A Tool for Determining Levels of Autonomy and Automation, Applied to Human-Rated Space Systems" in *Infotech@Aerospace*, American Institute of Aeronautics and Astronautics (AIAA) 2005-7061. Arlington, VA: AIAA.
- Proud, Ryan. W., and Jeremy J. Hart. 2005. "Function-specific Level of Autonomy and Automation Tool (FLOAAT) Rendezvous, Proximity Operations, and Docking (RPOD) Reference Levels of Autonomy and Automation" (NASA/JSC Document Number AFMFLOAAT002, Output Version). Unpublished.
- Psychology Dictionary* Web Site. 2014. "What Is Cooper-Harper Handling Qualities Rating Scale?" <http://psychologydictionary.org/cooper-harper-handling-qualities-rating-scale/>.
- Sheridan, Thomas B. 1987. "Supervisory Control." In *Handbook of Human Factors*, edited by Gavriel Salvendy, 1244–1268. New York: Wiley.

- Sheridan, Thomas B. 1992 *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, MA: MIT Press.
- Sheridan, Thomas. B. 2002. *Humans and Automation: System Design and Research Issues*. New York: Wiley.
- Sheridan, Thomas. B., and Gunnar A. Johannesen, eds. 1976. *Monitoring Behavior and Supervisory Control*. New York: Springer.
- Sheridan, Thomas B., and William L. Verplank. 1978. *Human and Computer Control Of Undersea Teleoperators*. Arlington, VA: Office of Naval Research.
- Sierhuis, Maarten, Jeffrey M. Bradshaw, Alessandro Acquisti, Ron van Hoof, Renia Jeffers, and Andrzej Uszok. 2003. "Human-Agent Teamwork and Adjustable Autonomy in Practice." Paper presented at the proceedings of the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), Nara, Japan, May 19–23, 2003.
- Sullivan, Don, Joseph Totah, Steve Wegener, Francis Enomoto, Chad Frost, John Kaneshige, and Jeremy Frank. 2004. *Intelligent Mission Management for Uninhabited Aerial Vehicles*. Moffett Field, CA: NASA Ames Research Center.
- Tecuci, Gheorghe, D. Aha, Mihai Boicu, Michael Cox, G. Ferguson, and A. Tate. 2003. "Workshop on Mixed-Initiative Intelligent Systems." Workshop presented at the proceedings of the 18th International Joint Conference on Artificial Intelligence, Acapulco, Mexico, August 9–15.
- UAS Control Segment (UCS) Working Group. 2014. *UAS Control Segment (UCS) Architecture Description Version 2.2 (UCS-INF-AD)*. Washington, DC, Office of the Secretary of Defense.
- Vaneman, Warren, and Kostos Triantis. 2014. "An Analytical Approach to Assessing Resilience in a System of Systems." PowerPoint presentation presented at the Systems Engineering in DC conference, Chantilly, Virginia, April 3–5.
- Zieba, Stéphane, David Jouglet, Philippe Polet, and Frédéric Vanderhaegen. 2007. "Resilience and Affordances: Perspectives for Human-Robot Cooperation?" Paper presented at the 26th European Annual Conference on Human Decision-Making and Manual Control (EAM'07), Copenhagen, Denmark, June 21–22.

Zieba, Stéphane, Philippe Polet, Frédéric Vanderhaegen, and Serge Debernard. 2009. "Resilience of a Human-Robot System Using Adjustable Autonomy and Human-Robot Collaborative Control." *International Journal of Adaptive and Innovative Systems* 1, no. 1: 13–29.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California